Network Working Group
Request for Comments: 2871
Category: Informational

J. Rosenberg
dynamicsoft
H. Schulzrinne
Columbia University

June 2000

A Framework for Telephony Routing over IP

Status of this Memo

Copyright Notice

Abstract

This document serves as a framework for Telephony Routing over IP (TRIP), which supports the discovery and exchange of IP telephony gateway routing tables between providers. The document defines the problem of telephony routing exchange, and motivates the need for the protocol. It presents an architectural framework for TRIP, defines terminology, specifies the various protocol elements and their functions, overviews the services provided by the protocol, and discusses how it fits into the broader context of Internet telephony.

Table of Contents

1 Introduction

This document serves as a framework for Telephony Routing over IP (TRIP), which supports the discovery and exchange of IP telephony gateway routing tables between providers. The document defines the problem of telephony routing exchange, and motivates the need for the protocol. It presents an architectural framework for TRIP, defines terminology, specifies the various protocol elements and their functions, overviews the services provided by the protocol, and discusses how it fits into the broader context of Internet telephony.

2 Terminology

We define the following terms. Note that there are other definitions for these terms, outside of the context of gateway location. Our definitions aren't general, but refer to the specific meaning here:

Gateway: A device with some sort of circuit switched network connectivity and IP connectivity, capable of initiating and terminating IP telephony signaling protocols, and capable of initiating and terminating telephone network signaling protocols.

End User: The end user is usually (but not necessarily) a human being, and is the party who is the ultimate initiator or recipient of calls.

Calling Device: The calling device is a physical entity which has IP connectivity. It is under the direction of an end user who wishes to place a call. The end user may or may not be directly controlling the calling device. If the calling device is a PC, the end user is directly controlling it. If, however, the calling device is a telephony gateway, the end user may be accessing it through a telephone.

Gatekeeper: The H.323 gatekeeper element, defined in [1].

SIP Server: The Session Initiation Protocol proxy or redirect server defined in [2].

Call Agent: The MGCP call agent, defined in [3].

GSTN: The Global Switched Telephone Network, which is the worldwide circuit switched network.

Signaling Server: A signaling server is an entity which is capable of receiving and sending signaling messages for some IP telephony signaling protocol, such as H.323 or SIP. Generally speaking, a signaling server is a gatekeeper, SIP server, or call agent.

Location Server (LS): A logical entity with IP connectivity which has knowledge of gateways that can be used to terminate calls towards the GSTN. The LS is the main entity that participates in Telephony Routing over IP. The LS is generally a point of contact for end users for completing calls to the telephony network. An LS may also be responsible for propagation of gateway information to other LS’s. An LS may be coresident with an H.323 gatekeeper or SIP server, but this is not required.

Internet Telephony Administrative Domain (ITAD): The set of resources (gateways and Location Servers) under the control of a single administrative authority. End users are customers of an ITAD.

Provider: The administrator of an ITAD.

Location Server Policy: The set of rules which dictate how a location server processes information it sends and receives via TRIP. This includes rules for aggregating, propagating, generating, and accepting information.

End User Policy: Preferences that an end user has about how a call towards the GSTN should be routed.

Peers: Two LS’s are peers when they have a persistent association between them over which gateway information is exchanged.

Internal peers: Peers that both reside within the same ITAD.

External peers: Peers that reside within different ITADs.

Originating Location Server: A Location Server which first generates a route to a gateway in its ITAD.

Telephony Routing Information Base (TRIB): The database of gateways an LS builds up as a result of participation in TRIP.

3 Motivation and Problem Definition

As IP telephony gateways grow in terms of numbers and usage, managing their operation will become increasingly complex. One of the difficult tasks is that of gateway location, also known as gateway selection, path selection, gateway discovery, and gateway routing. The problem occurs when a calling device (such as a telephony gateway or a PC with IP telephony software) on an IP network needs to complete a call to a phone number that represents a terminal on a circuit switched telephone network. Since the intended target of the call resides in a circuit switched network, and the caller is initiating the call from an IP host, a telephony gateway must be used. The gateway functions as a conversion point for media and signaling, converting between the protocols used on the IP network, and those used in the circuit switched network.

The gateway is, in essence, a relaying point for an application layer signaling protocol. There may be many gateways which could possibly complete the call from the calling device on the IP network to the called party on the circuit switched network. Choosing such a gateway is a non-trivial process. It is complicated because of the following issues:

Number of Candidate Gateways: It is anticipated that as IP telephony becomes widely deployed, the number of telephony gateways connecting the Internet to the GSTN will become large. Attachment to the GSTN means that the gateway will have connectivity to the nearly one billion terminals reachable on this network. This means that every gateway could theoretically complete a call to any terminal on the GSTN. As such, the number of candidate gateways for completing a call may be very large.

Business Relationships: In reality, the owner of a gateway is unlikely to make the gateway available to any user who wishes to connect to it. The gateway provides a useful service, and incurs cost when completing calls towards the circuit switched network. As a result, providers of gateways will, in many cases, wish to charge for use of these gateways. This may restrict usage of the gateway to those users who have, in some fashion, an established relationship with the gateway provider.

Provider Policy: In all likelihood, an end user who wishes to make use of a gateway service will not compensate the gateway provider directly. The end user may have a relationship with an IP telephony service provider which acts as an intermediary to providers of gateways. The IP telephony service provider may have gateways of its own as well. In this case, the IP telephony service provider may have policies regarding the usage of various gateways from other providers by its customers. These policies must figure into the selection process.

End User Policy: In some cases, the end user may have specific requirements regarding the gateway selection. The end user may need a specific feature, or have a preference for a certain provider. These need to be taken into account as well.

Capacity: All gateways are not created equal. Some are large, capable of supporting hundreds or even thousands of simultaneous calls. Others, such as residential gateways, may only support one or two calls. The process for selecting gateways should allow gateway capacity to play a role. It is particularly desirable to support some form of load balancing across gateways based on their capacities.

Protocol and Feature Compatibilities: The calling party may be using a specific signaling or media protocol that is not supported by all gateways.

From these issues, it becomes evident that the selection of a gateway is driven in large part by the policies of various parties, and by the relationships established between these parties. As such, there cannot be a global "directory of gateways" in which users look up phone numbers. Rather, information on availability of gateways must be exchanged by providers, and subject to policy, made available locally and then propagated to other providers. This would allow each provider to build up its own local database of available gateways - such a database being very different for each provider depending on policy.

From this, we can conclude that a protocol is needed between administrative domains for exchange of gateway routing information. The protocol that provides these functions is Telephony Routing over IP (TRIP). TRIP provides a specific set of functions:

- o Establishment and maintenance of peering relationships between providers;
- o Exchange and synchronization of telephony gateway routing information between providers;
- o Prevention of stable routing loops for IP telephony signaling protocols;
- o Propagation of learned gateway routing information to other providers in a timely and scalable fashion;
- o Definition of the syntax and semantics of the data which describe telephony gateway routes.

TRIP can be generally summarized as an inter-domain IP telephony gateway routing protocol.

4 Related Problems

At a high level, the problem TRIP solves appears to be a mapping problem: given an input telephone number, determine, based on some criteria, the address of a telephony gateway. For this reason, the gateway location problem is often called a "phone number to IP address translation problem". This is an over-simplification, however. There are at least three separate problems, all of which can be classified as a "phone number to IP address translation problem", and only one of which is addressed by TRIP:

- o Given a phone number that corresponds to a terminal on a circuit switched network, determine the IP address of a gateway capable of completing a call to that phone number.
- o Given a phone number that corresponds to a specific host on the Internet (this host may have a phone number in order to facilitate calls to it from the circuit switched network), determine the IP address of this host.
- o Given a phone number that corresponds to a user of a terminal on a circuit switched network, determine the IP address of an IP terminal which is owned by the same user.

The last of these three mapping functions is useful for services where the PC serves as an interface for the phone. One such service is the delivery of an instant message to a PC when the user's phone rings. To deliver this service, a switch in the GSTN is routing a call towards a phone number. It wishes to send an Instant Message to the PC for this user. This switch must somehow have access to the IP network, in order to determine the IP address of the PC corresponding to the user with the given phone number. The mapping function is a name to address translation problem, where the name happens to be represented by a string of digits. Such a translation function is best supported by directory protocols. This problem is not addressed by TRIP.

The second of these mappings is needed to facilitate calls from traditional phones to IP terminals. When a user on the GSTN wishes to call a user with a terminal on the IP network, they need to dial a number identifying that terminal. This number could be an IP address. However, IP addresses are often ephemeral, assigned on demand by DHCP [4] or by dialup network access servers using PPP [5]. The number could be a hostname, obtained through some translation of groups of numbers to letters. However, this is cumbersome. It has been proposed instead to assign phone numbers to IP telephony terminals. A caller on the GSTN would then dial this number as they would any other. This number serves as an alternate name for the IP terminal, in much the same way its hostname serves as a name. A switch in the GSTN must then access the IP network, and obtain the mapping from this number to an IP address for the PC. Like the previous case, this problem is a name to address translation problem, and is best handled by a directory protocol. It is not addressed by TRIP.

The first mapping function, however, is fundamentally an address to route translation problem. It is this problem which is considered by TRIP. As discussed in Section 3, this mapping depends on local factors such as policies and provider relationships. As a result, the database of available gateways is substantially different for each provider, and needs to be built up through specific inter-provider relationships. It is for this reason that a directory protocol is not appropriate for TRIP, whereas it is appropriate for the others.

5 Relationship with BGP

TRIP can be classified as a close cousin of inter-domain IP routing protocols, such as BGP [6]. However, there are important differences between BGP and TRIP:

- TRIP runs at the application layer, not the network layer, where BGP resides.

- TRIP runs between servers which may be separated by many intermediate networks and IP service providers. BGP runs between routers that are usually adjacent.

- o The information exchanged between TRIP peers describes routes to application layer devices, not IP routers, as is done with BGP.

- o TRIP assumes the existence of an underlying IP transport network. This means that servers which exchange TRIP routing information need not act as forwarders of signaling messages that are routed based on this information. This is not true in BGP, where the peers must also act as forwarding points (or name an adjacent forwarding hop) for IP packets.

- o The purpose of TRIP is not to establish global connectivity across all ITADs. It is perfectly reasonable for there to be many small islands of TRIP connectivity. Each island represents a closed set of administrative relationships. Furthermore, each island can still have complete connectivity to the entire GSTN. This is in sharp contrast to BGP, where the goal is complete connectivity across the Internet. If a set of AS's are isolated from some other set because of a BGP disconnect, no IP network connectivity exists between them.

- o Gateway routes are far more complex than IP routes (since they reside at the application, not the network layer), with many more parameters which may describe them.

- o BGP exchanges prefixes which represent a portion of the IP name space. TRIP exchanges phone number ranges, representing a portion of the GSTN numbering space. The organization and hierarchies in these two namespaces are different.

These differences means that TRIP borrows many of the concepts from BGP, but that it is still a different protocol with its own specific set of functions.

6 Example Applications of TRIP

TRIP is a general purpose tool for exchanging IP telephony routes between providers. TRIP does not, in any way, dictate the structure or nature of the relationships between those providers. As a result, TRIP has applications for a number of common cases for IP telephony.

6.1 Clearinghouses

A clearinghouse is a provider that serves as an exchange point between a number of other providers, called the members of the clearinghouse. Each member signs on with the clearinghouse. As part of the agreement, the member makes their gateways available to the other members of the clearinghouse. In exchange, the members have access to the gateways owned by the other members of the clearinghouse. When a gateway belonging to one member makes a call, the clearinghouse plays a key role in determining which member terminates the call.

TRIP can be applied here as the tool for exchanging routes between the members and the clearinghouse. This is shown in Figure 1.

There are 6 member companies, M1 through M6. Each uses TRIP to send and receive gateway routes with the clearinghouse provider.

6.2 Confederations

We refer to a confederation as a group of providers which all agree to share gateways with each other in a full mesh, without using a central clearinghouse. Such a configuration is shown in Figure 2. TRIP would run between each pair of providers.

6.3 Gateway Wholesalers

```
          ------                                  ------
         |      |                                |      |
         | M1   |    TRIP                 TRIP   |  M2  |
         |      |\    |                    |    /|      |
          ------  \   |                    |   /  ------
                   \ \ /   -------------- \ / /
          ------    \----|              |----/    ------
         |      |        |              |        |      |
         | M3   |--------| Clearinghouse|--------|  M4  |
         |      |        |              |        |      |
          ------    /----|              |----\    ------
                   /      --------------      \
          ------  /                            \  ------
         |      |/                              \|      |
         | M5   |                                |  M6  |
         |      |                                |      |
          ------                                  ------
```

Figure 1: TRIP in the Clearinghouse Application

```
                   ------        ------
                  |      |------|      |
                  | M1   |      | M2   |
                  |      |\    /|      |
                   ------  \  /  ------
                     |      \/      |
                     |      /\      |<-----TRIP
                   ------  /  \  ------
                  |      |/    \|      |
                  | M3   |      | M4   |
                  |      |------|      |
                   ------        ------
```

Figure 2: TRIP for Confederations

In this application, there are a number of large providers of telephony gateways. Each of these resells its gateway services to medium sized providers. These, in turn, resell to local providers who sell directly to consumers. This is effectively a pyramidal relationship, as shown in Figure 3.

```
                         ------
                        |      |
                        | M1   |
                        |      |
                         ------
                       /        \ <------- TRIP
                 ------        ------
                |      |      |      |
                | M2   |      | M3   |
                |      |      |      |
                 ------        ------
                /      \      /      \
          ------        ------        ------
         |      |      |      |      |      |
         | M4   |      | M5   |      | M6   |
         |      |      |      |      |      |
          ------        ------        ------
```

Figure 3: TRIP for Wholesalers

Note that in this example, provider M5 resells gateways from both M2 and M3.

7 Architecture

Figure 4 gives the overall architecture of TRIP.

```
        ITAD1                                ITAD2
   -----------------                 ------------------
  |                   |             |                   |
  |  ----             |             |            ----   |
  | | GW |            |             |           | EU |  |
  |  ----  \   ----   |             |   ----  /  ----   |
  |           | LS | ---------------- | LS |            |
  |  ----      ----   |            /   ----  \   ----   |
  | | GW | /          |           /|            | EU |  |
  |  ----             |          / |             ----   |
  |                   |         /  |                    |
   ------------------          /    -----------------
                              /
                             /
                 --------- /----------
                |         |            |
                |        ----          |
                |       | LS |         |
                |    /   ---- \        |
                |  ----   ||    ----   |
                | | GW |  ||   | EU |  |
                |  ----   ||    ----   |
                |  ----   ||    ----   |
                | | GW | /  \  | EU |  |
                |  ----         ----   |
                |                      |
                 ---------------------
                         ITAD3
```

Figure 4: TRIP Architecture

There are a number of Internet Telephony administrative domains (ITAD's), each of which has at least one Location Server (LS). The LS's, through an out-of-band means, called the intra-domain protocol, learn about the gateways in their domain. The intra-domain protocol is represented by the lines between the GW and LS elements in ITAD1 in the Figure. The LS's have associations with other LS's, over which they exchange gateway information. These associations are established administratively, and are set up when the IT administrative domains have some kind of agreements in place regarding exchange of gateway information. In the figure, the LS in ITAD1 is connected to the LS in ITAD2, which is in turn connected to the LS in ITAD3. Through Telephony Routing over IP (TRIP), the LS in ITAD2 learns about the two gateways in ITAD1. This information is accessed by end users

(EUs) in ITAD2 through the front-end. The front-end is a non-TRIP protocol or mechanism by which the LS databases are accessed. In ITAD3, there are both EUs and gateways. The LS in ITAD3 learns about the gateways in ITAD1 through a potentially aggregated advertisement from the LS in ITAD2.

8 Elements

The architecture in Figure 4 consists of a number of elements. These include the IT administrative domain, end user, gateway, and location server.

8.1 IT Administrative Domain

An IT administrative domain consists of zero or more gateways, at least one Location Server, and zero or more end users. The gateways and LS's are those which are under the administrative control of a single authority. This means that there is one authority responsible for dictating the policies and configuration of the gateways and LS's.

An IT administrative domain need not be the same as an autonomous system. While an AS represents a set of physically connected networks, an IT administrative domain may consist of elements on disparate networks, and even within disparate autonomous systems.

The end users within an IT administrative domain are effectively the customers of that IT administrative domain. They are interested in completing calls towards the telephone network, and thus need access to gateways. An end user may be a customer of one IT administrative domain for one call, and then a customer of a different one for the next call.

An IT administrative domain need not have any gateways. In this case, its LS learns about gateways in other domains, and makes these available to the end users within its domain. In this case, the IT administrative domain is effectively a virtual IP telephony gateway provider. This is because it provides gateway service, but may not actually own or administer any gateways.

An IT administrative domain need not have any end users. In this case, it provides "wholesale" gateway service, making its gateways available to customers in other IT administrative domains.

An IT administrative domain need not have gateways nor end users. In this case, the ITAD only has LS's. The ITAD acts as a reseller, learning about other gateways, and then aggregating and propagating this information to other ITAD's which do have customers.

8.2 Gateway

A gateway is a logical device which has both IP connectivity and connectivity to some other network, usually a public or private telephone network. The function of the gateway is to translate the media and signaling protocols from one network technology to the other, achieving a transparent connection for the users of the system.

A gateway has a number of attributes which characterize the service it provides. Most fundamental among these are the range of phone numbers to which it is willing to provide service. This range may be broken into subranges, and associated with each, some cost metric or cost token. This token indicates some notion of cost or preference for completing calls for this part of the telephone number range.

A gateway has attributes which characterize the volume of service which it can provide. These include the number of ports it has (i.e., the number of simultaneous phone calls it can support), and the access link speed. These two together represent some notion of the capacity of the gateway. The metric is useful for allowing Location Servers to decide to route calls to gateways in proportion to the value of the metric, thus achieving a simple form of load balancing.

A gateway also has attributes which characterize the type of service it provides. This includes, but is not limited to, signaling protocols supported, telephony features provided, speech codecs understood, and encryption algorithms which are implemented. These attributes may be important in selecting a gateway. In the absence of baseline required features across all gateways (an admirable, but difficult goal), such a set of attributes is required in order to select a gateway with which communications can be established. End users which have specific requirements for the call (such as a user requesting a business class call, in which case certain call features may need to be supported) may wish to make use of such information as well.

Some of these attributes are transported in TRIP to describe gateways, and others are not. This depends on whether the metric can be reasonably aggregated, and whether it is something which must be conveyed in TRIP before the call is set up (as opposed to negotiated or exchanged by the signaling protocols themselves). The philosophy of TRIP is to keep it simple, and to favor scalability above abundance of information. TRIP's attribute set is readily extensible. Flags provide information that allow unknown attributes to be reasonably processed by an LS.

8.3 End Users

An end user is an entity (usually a human being) which wishes to complete a call through a gateway from an IP network to a terminal on a telephone network. An end user may be a user logged on at a PC with some Internet telephony software. The end user may also be connected to the IP network through an ingress telephone gateway, which the user accessed from telephone handset. This is the case for what is referred to as "phone to phone" service with the IP network used for interexchange transport.

End users may, or may not be aware that there is a telephony routing service running when they complete a call towards the telephone network. In cases where they are aware, end users may have preferences for how a call is completed. These preferences might include call features which must be supported, quality metrics, owner or administrator, and cost preferences.

TRIP does not dictate how these preferences are combined with those of the provider to yield the final gateway selection. Nor does TRIP support the transport of these preferences to the LS. This transport can be accomplished using the front end, or by some non-protocol means.

8.4 Location Server

The Location Server (LS) is the main functional entity of TRIP. It is a logical device which has access to a database of gateways, called the Telephony Routing Information Base (TRIB). This database of gateways is constructed by combining the set of locally available gateways and the set of remote gateways (learned through TRIP) based on policy. The LS also exports a set of gateways to its peer LS's in other ITAD's. The set of exported gateways is constructed from the set of local gateways and the set of remote gateways (learned through TRIP) based on policy. As such, policy plays a central role in the LS operation. This flow of information is shown in Figure 5.

```
                              |
                              |Intra-domain protocol
                             \ /
                            Local
                           Gateways


TRIP-->  Gateways    POLICY     Gateways -->TRIP
             IN                       Out
                                |
                               \ /
                          Telephony Routing
                          Information Base
```

Figure 5: Flow of Information in TRIP

The TRIB built up in the LS allows it to make decisions about IP telephony call routing. When a signaling message arrives at a signaling server, destined for a telephone network address, the LS's database can provide information which is useful for determining a gateway or an additional signaling server to forward the signaling message to. For this reason, an LS may be coresident with a signaling server. When they are not coresident, some means of communication between the LS and the signaling server is needed. This communication is not specifically addressed by TRIP, although it is possible that TRIP might meet the needs of such a protocol.

An ITAD must have at least one LS in order to participate in TRIP. An ITAD may have more than one LS, for purposes of load balancing, ease of management, or any other reason. In that case, communications between these LS's may need to take place in order to synchronize databases and share information learned from external peers. This is often referred to as the interior component of an inter-domain protocol. TRIP includes such a function.

Figure 5 shows an LS learning about gateways within the ITAD by means of an intra-domain protocol. There need not be an intra-domain protocol. An LS may operate without knowledge of any locally run gateways. Or, it may know of locally run gateways, but through static configuration. An LS may also be co-resident with a gateway, in which case it would know about the gateway that it is co-resident with.

9 Element Interactions

9.1 Gateways and Location Servers

Gateways must somehow propagate information about their characteristics to an LS within the same ITAD. This LS may, in turn, further propagate this information outside of the ITAD by means of TRIP. This LS is called an originating LS for that gateway. When an LS nis not coresident with the gateway, the means by which the information gets propagated is not within the scope of TRIP. The protocol used to accomplish this is generally called an intra-domain protocol.

One way in which the information can be propagated is with the Service Location Protocol (SLP) [7]. The gateway can contain a Service Agent (SA), and the LS can act as a Directory Agent (DA). SLP defines procedures by which service information is automatically propagated to DA's from SA's. In this fashion, an LS can learn about gateways in the ITAD.

An alternate mechanism for the intra-domain protocol is via the registration procedures of SIP or H.323. The registration procedures provide a means by which users inform a gatekeeper or SIP server about their address. Such a registration procedure could be extended to allow a gateway to effectively register as well.

LDAP [8] might also be used for the intra-domain protocol. A gateway can use LDAP to add an entry for itself into the database. If the LS also plays the role of the LDAP server, it will be able to learn about all those gateways in its ITAD.

The intra-domain protocol which is used may be different from IT administrative domain to IT administrative domain, and is a matter of local configuration. There may also be more than one intra-domain protocol in a particular ITAD. An LS can also function without an intra-domain protocol. It may learn about gateways through static configuration, or may not know of any local gateways.

9.2 Location Server to Location Server

The interaction between LS's is what is defined by TRIP. LS's within the same ITAD use TRIP to synchronize information amongst themselves. LS's within different ITADs use TRIP to exchange gateway information according to policy. In the former case the LS's are referred to as internal peers, and in the latter case, external peers.

LS’s communicate with each other through persistent associations. An LS may be connected to one or more other LS’s. LS’s need not be physically adjacent or part of the same autonomous system. The association between a pair of LS’s is normally set up administratively. Two LS’s are configured to communicate with each other when their administrators have an agreement in place to exchange gateway information. While TRIP does not provide an autodiscovery procedure for peer LS’s to discover each other, one could possibly be used. Such a procedure might be useful for finding a backup peer LS when a crash occurs. Alternatively, in an environment where the business relationships between peers become more standardized, peers might be allowed to discover each other through protocols like the Service Location Protocol (SLP) [9]. Determination about whether autodiscovery should or should not be used is at the discretion of the administrator.

The syntax and semantics of the messages exchanged over the association between LS’s are dictated by TRIP. The protocol does not dictate the nature of the agreements which must be in place. TRIP merely provides a transport means to exchange whatever gateway routing information is deemed appropriate by the administrators of the system. Details are provided in the TRIP protocol specification itself.

The rules which govern which gateway information is generated, propagated, and accepted by a gateway is called a location server policy. TRIP does not dictate or mandate any specific policy.

9.2.1 Nature of Exchanged Information

The information exchanged by the LS’s is a set of routing objects. Each routing object minimally consists of a range of telephone numbers which are reachable, and an IP address or host name which is the application-layer "next hop" towards a gateway which can reach that range. Routing objects are learned from the intra-domain protocol, static configuration, or from LS’s in remote ITAD’s. An LS may aggregate these routing objects together (merging ranges of telephone numbers, and replacing the IP address with its own IP address, or with the IP address of a signaling server with which the LS is communicating) and then propagate them to another LS. The decision about which objects to aggregate and propagate is known as a route selection operation. The administrator has great latitude in selecting which objects to aggregate and propagate, so long as they are within the bounds of correct protocol operation (i.e., no loops are formed). The selection can be made based on information learned through TRIP, or through any out of band means.

A routing object may have additional information which characterizes the service at the gateway. These attributes include things like protocols, features supported, and capacity. Greater numbers of attributes can provide useful information, however, they come at a cost. Aggregation becomes difficult with more and more information, impacting the scalability of the protocol.

Aggregation plays a central role in TRIP. In order to facilitate scalability, routing objects can be combined into larger aggregates before being propagated. The mechanisms by which this is done are specified in TRIP. Aggregation of application layer routes to gateways is a non-trivial problem. There is a fundamental tradeoff between aggregatability and verbosity. The more information that is present in a TRIP routing object, the more difficult it is to aggregate.

Consider a simple example of two gateways, A and B, capable of reaching some set of telephone numbers, X and Y, respectively. C is an LS for the ITAD in which A and B are resident. C learns of A and B through some other means. As it turns out, X and Y can be combined into a single address range, Z. C has several options. It can propagate just the advertisement for A, just the advertisement for B, propagate both, or combine them and propagate the aggregate advertisement. In this case C chooses the latter approach, and sends a single routing object to one of its peers, D, containing address range Z and its own address, since it is also a signaling server. D is also a signaling server.

Some calling device, E, wishes to place a phone call to telephone number T, which happens to be in the address range X. E is configured to use D as its default H.323 gatekeeper. So, E sends a call setup message to D, containing destination address T. D determines that the address T is within the range Z. As D had received a routing object from C containing address range Z, it forwards the call setup message to C. C, in turn, sees that T is within range X, and so it forwards the call setup to A, which terminates the call signaling and initiates a call towards the telephone network.

9.2.2 Quality of Service

One of the factors which is useful to consider when selecting a gateway is "QoS" - will a call through this gateway suffer sufficiently low loss, delay, and jitter? The quality of a call depends on two components - the QoS on the path between the caller and gateway, and the capacity of the gateway itself (measured in terms of number of circuits available, link capacity, DSP resources, etc.). Determination of the latter requires intricate knowledge of underlying network topologies, and of where the caller is located. This is something handled by QoS routing protocols, and is outside the scope of TRIP.

However, gateway capacity is not dependent on the caller location or path characteristics. For this reason, a capacity metric of some form is supported by TRIP. This metric represents the static capacity of the gateway, not the dynamic available capacity which varies continuously during the gateways operation. LS's can use this metric as a means of load balancing of calls among gateways. It can also be used as an input to any other policy decision.

9.2.3 Cost Information

Another useful attribute to propagate is a pricing metric. This might represent the amount a particular gateway might charge for a call. The metric can be an index into a table that defines a pricing structure according to a pre-existing business arrangement, or it can contain a representation of the price itself. TRIP itself does not define a pricing metric, but one can and should be defined as an extension. Using an extension for pricing means more than one such metric can be defined.

10 The Front End

As a result of TRIP, the LS builds up a database (the TRIB) of gateway routes. This information is made available to various entities within the ITAD. The way in which this information is made available is called the front end. It is the visible means by which TRIP services are exposed outside of the protocol.

10.1 Front End Customers

There are several entities which might use the front end to access the TRIB. These include, but are not limited to:

Signaling Servers: Signaling servers receive signaling messages (such as H.323 or SIP messages) whose purpose is the initiation of IP telephony calls. The destination address of these calls may be a phone number corresponding to a terminal on the GSTN. In order to route these calls to an appropriate gateway, the signaling server will need access to the database built up in the LS.

End Users: End users can directly query the LS to get routing information. This allows them to provide detailed information on their requirements. They can then go and contact the next hop signaling server or gateway towards that phone number.

Administrators: Administrators may need to access the TRIB for maintenance and management functions.

When a signaling server contacts the LS to route a phone number, it is usually doing so because a calling device (on behalf of an end user) has attempted to set up a call. As a result, signaling servers effectively act as proxies for end users when accessing the LS database. The communication between the calling devices and their proxies (the signaling servers) is through the signaling protocol.

The advantage of this proxy approach is that the actual LS interaction is hidden from the calling device. Therefore, whether the call is to a phone number or IP address is irrelevant. The routing in the case of phone numbers takes place transparently. Proxy mode is also advantageous for thin clients (such as standalone IP telephones) which do not have the interfaces or processing power for a direct query of the LS.

The disadvantage of the proxy approach is the same as its advantage - the LS interaction is hidden from the calling device (and thus the end user). In some cases, the end user may have requirements as to how they would like the call to be routed. These include preferences about cost, quality, administrator, or call services and protocols. These requirements are called the end user policy. In the proxy approach, the user effectively accesses the service through the signaling protocol. The signaling protocol is not likely to be able to support expression of complex call routing preferences from end users (note however, that SIP does support some forms of caller preferences for call routing [10]). Therefore, direct access from the end user to the LS can provide much richer call routing services.

When the end user policy is presented to the LS (either directly or through the signaling protocol), it is at the discretion of the LS how to make use of it. The location server may have its own policies regarding how end user preferences are handled.

10.2 Front End Protocols

There are numerous protocols that can be used in the front end to access the LS database. TRIP does not specify or restrict the possibilities for the front end. It is not clear that it is necessary or even desirable for there to be a single standard for the front end. The various protocols have their strengths and weaknesses. One may be the right solution in some cases, and another in different cases.

Some of the possible protocols for the front end are:

Service Location Protocol (SLP): SLP has been designed to fit exactly this kind of function. SLP is ideal for locating servers described by a set of attributes. In this case, the server is a gateway (or next hop towards the gateway), and the attributes are the end user policy. The end user is an SLP UA, and the LS is an SLP DA. The Service Query is used to ask for a gateway with a particular set of attributes.

Open Settlements Protocol (OSP): OSP [11] is a client server protocol. It allows the client to query a server with a phone number, and get back the address of a next hop, along with authorization tokens to use for the call. In this case, the server can be an LS. The routing table it uses to respond to OSP queries is the one built up using TRIP.

Lightweight Directory Access Protocol (LDAP): LDAP is used for accessing distributed databases. Since the LS server contains a database, LDAP could be used to query it.

Web Page: The LS could have a web front end. Users could enter queries into a form, and the matching gateways returned in the response. This access mechanism is more appropriate for human access, however. A signaling server would not likely access the front end through a web page.

TRIP: The protocols discussed above are all of the query-response type. There is no reason why the LS access must be of this form. It is perfectly acceptable for the access to be through complete database synchronization, so that the entity accessing the LS database effectively has a full copy of it. If this approach were desired, TRIP itself is an appropriate mechanism. This approach has obvious drawbacks, but nothing precludes it from being done.

11 Number Translations

The model for TRIP is that of many gateways, each of which is willing to terminate calls towards some set of phone numbers. Often, this set will be based on the set of telephone numbers which are in close geographic proximity to the gateway. For example, a gateway in New York might be willing to terminate calls to the 212 and 718 area codes. Of course, it is up to the administrator to decide on what phone numbers the gateway is willing to call.

However, certain phone numbers don't represent GSTN terminals at all, but rather they represent services or virtual addresses. An example of such numbers are freephone and LNP numbers. In the telephone network, these are actually mapped to routable telephone numbers, often based on complex formulae. A classic example is time-of-day-based translation.

While nothing prevents a gateway from advertising reachability to these kinds of numbers, this usage is highly discouraged. Since TRIP is a routing protocol, the routes it propagates should be to routable numbers, not to names which are eventually translated to routable numbers. Numerous problems arise when TRIP is used to propagate routes to these numbers:

- o Often, these numbers have only local significance. Calls to a freephone number made from New York might terminate in a New York office of a company, while calls made from California will terminate in a California branch. If this freephone number is injected into TRIP by a gateway in New York, it could be propagated to other LS's with end users in California. If this route is used, calls may be not be routed as intended.

- o The call signaling paths might be very suboptimal. Consider a gateway in New York that advertises a ported number that maps to a phone in California. This number is propagated by TRIP, eventually being learned by an LS with end users in California. When one of them dials this number, the call is routed over the IP network towards New York, where it hits the gateway, and then is routed over the GSTN back to California. This is a waste of resources. Had the ported number been translated before the gateway routing function was invoked, a California gateway could have been accessed directly.

As a result, it is more efficient to perform translations of these special numbers before the LS routing databases are accessed. How this translation is done is outside the scope of TRIP. It can be accomplished by the calling device before making the call, or by a signaling server before it accesses the LS database.

12 Security Considerations

Security is an important component in TRIP. The TRIP model assumes a level of trust between peer LS's that exchange information. This information is used to propagate information which determines where calls will be routed. If this information were incorrect, it could cause complete misrouting of calls. This enables a significant denial of service attack. The information might also be propagated to other

ITADs, causing the problem to potentially spread. As a result, mutual authentication of peer LS’s is critical. Furthermore, message integrity is required.

TRIP messages may contain potentially sensitive information. They represent the routing capabilities of an ITAD. Such information might be used by corporate competitors to determine the network topology and capacity of the ITAD. As a result, encryption of messages is also supported in TRIP.

As routing objects can be passed via one LS to another, there is a need for some sort of end to end authentication as well. However, aggregation will cause the routing objects to be modified, and therefore authentication can only take place from the point of last aggregation to the receiving LS’s.

13 Acknowledgments

The authors would like to thank Randy Bush, Mark Foster, Dave Oran, Hussein Salama, and Matt Squire for their useful comments on this document.

14 Bibliography

[1] International Telecommunication Union, "Visual telephone systems and equipment for local area networks which provide a non-guaranteed quality of service," Recommendation H.323, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, May 1996.

[2] Handley, M., Schulzrinne, H., Schooler, E. and J. Rosenberg, "SIP: Session Initiation Protocol", RFC 2543, March 1999.

[3] Arango, M., Dugan, A., Elliott, I., Huitema, C. and S. Pickett, "Media Gateway Control Protocol (MGCP) Version 1.0", RFC 2705, October 1999.

[4] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[5] Simpson, W., "The Point-to-Point Protocol (PPP)," STD 51, RFC 1661, July 1994.

[6] Rekhter Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.

[7] Veizades, J., Guttman, E., Perkins, C. and S. Kaplan, "Service Location Protocol", RFC 2165, June 1997.

[8] Yeong, W., Howes, T. and S. Kille, "Lightweight Directory Access Protocol", RFC 1777, March 1995.

[9] Guttman, E., Perkins, C., Veizades, J. and M. Day, "Service Location Protocol, Version 2", RFC 2608, June 1999.

[10] Schulzrinne H. and J. Rosenberg, "SIP caller preferences and callee capabilities", Work in progress.

[11] European Telecommunications Standards Institute (ETSI), Telecommunications and Internet Protocol Harmonization Over Networks (TIPHON), "Inter-domain pricing, authorization, and usage exchange," Technical Specification 101 321 version 1.4.2, ETSI, 1998.

15 Authors' Addresses

Jonathan Rosenberg
dynamicsoft
72 Eagle Rock Avenue
First Floor
East Hanover, NJ 07936

Email: jdrosen@dynamicsoft.com

Henning Schulzrinne
Columbia University
M/S 0401
1214 Amsterdam Ave.
New York, NY 10027-7003

Email: schulzrinne@cs.columbia.edu

16. Full Copyright Statement

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.