

Network Working Group
Request for Comments: 5140
Category: Standards Track

M. Bangalore
R. Kumar
J. Rosenberg
Cisco
H. Salama
Citex Software
D.N. Shah
Moowee Inc.
March 2008

A Telephony Gateway REgistration Protocol (TGREP)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document describes the Telephony Gateway Registration Protocol (TGREP) for registration of telephony prefixes supported by telephony gateways and soft switches. The registration mechanism can also be used to export resource information. The prefix and resource information can then be passed on to a Telephony Routing over IP (TRIP) Location Server, which in turn can propagate that routing information within and between Internet Telephony Administrative Domains (ITADs). TGREP shares a lot of similarities with the TRIP protocol. It has similar procedures and finite state machine for session establishment. It also shares the same format for messages and a subset of attributes with TRIP.

Table of Contents

1. Introduction	3
2. Terminology and Definitions	5
3. TGREP: Overview of Operation	6
4. TGREP Attributes	7
4.1. TotalCircuitCapacity Attribute	8
4.2. AvailableCircuits Attribute	9
4.3. CallSuccess Attribute	10
4.4. Prefix Attributes	12
4.5. TrunkGroup Attribute	13
4.6. Carrier Attribute	15
5. TrunkGroup and Carrier Address Families	16
5.1. Address Family Syntax	17
6. Gateway Operation	18
6.1. Session Establishment	18
6.2. UPDATE Messages	19
6.3. KEEPALIVE Messages	19
6.4. Error Handling and NOTIFICATION Messages	19
6.5. TGREP Finite State Machine	19
6.6. Call Routing Databases	19
6.7. Multiple Address Families	20
6.8. Route Selection and Aggregation	20
7. LS/Proxy Behavior	20
7.1. Route Consolidation	22
7.2. Aggregation	23
7.3. Consolidation versus Aggregation	23
8. Security Considerations	23
9. IANA Considerations	24
9.1. Attribute Codes	24
9.2. Address Family Codes	24
10. Acknowledgments	25
11. References	25
11.1. Normative References	25
11.2. Informative References	26

1. Introduction

It is assumed that the reader of this document is familiar with TRIP [2, 12]. In typical Voice over IP (VoIP) networks, Internet Telephony Administrative Domains (ITADs) will deploy numerous gateways for the purposes of geographical diversity, capacity, and redundancy. When a call arrives at the domain, it must be routed to one of those gateways. Frequently, an ITAD will break its network into geographic Points of Presence (POPs), with each POP containing some number of gateways, and a proxy server element that fronts those gateways. The proxy element is a SIP proxy [11] or H.323 gatekeeper. The proxy server is responsible for managing the access to the POP, and also for determining which of the gateways will receive any given call that arrives at the POP. In conjunction with the proxy server that routes the call signaling, there are two components, the "Ingress LS" (aka "TGREP receiver") and the "Egress LS". The TGREP receiver component maintains TGREP peering relationship with one or more gateways. The routing information received from the gateways is further injected into the Egress LS, which in turn disseminates into the rest of the network on TRIP. For convenience, gateway and GW are used interchangeably.

This configuration is depicted graphically in Figure 1.

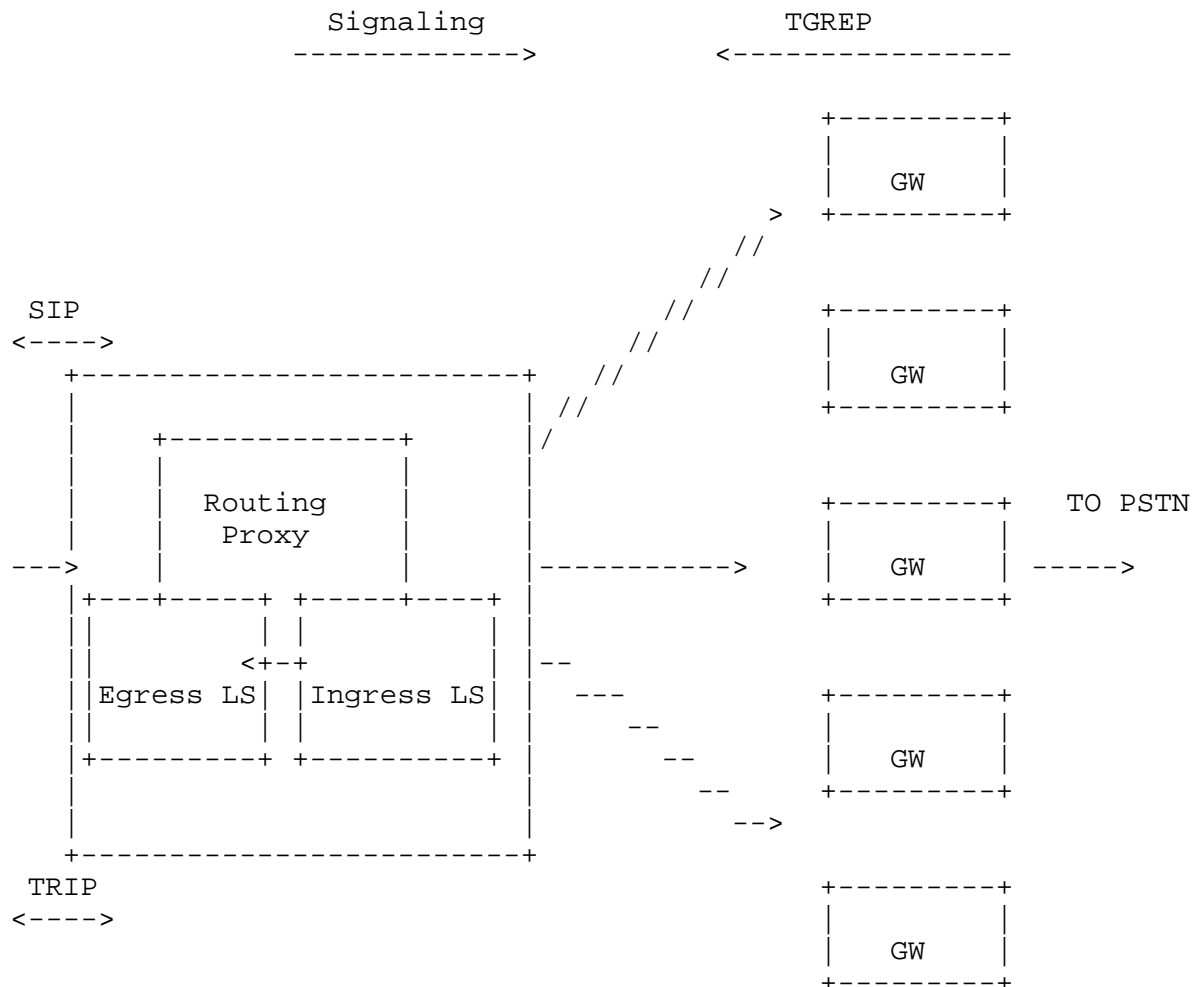


Figure 1: Gateway and LS Configuration

The decision about which gateway to use depends on many factors, including their availability, remaining call capacity, and call success statistics to a particular Public Switched Telephone Network (PSTN) destination (see [14]). For the proxy to do this adequately, it needs to have access to this information in real-time, as it changes. This means there must be some kind of communications between the proxy and the gateways to convey this information.

The TRIP protocol [2] is defined for carrying telephony routing information between providers, for the purposes of getting a call routed to the right provider for termination to the PSTN. However, there is no mechanism defined in TRIP that defines how routes get injected into the TRIP protocol from within the network. Nor does it

define mechanisms that would allow the provider to select the specific gateway for terminating a call when it arrives. Those gaps are filled by TGREP.

TGREP allows PSTN gateways or soft switches to inform a signaling server, such as a SIP proxy server or H.323 gatekeeper, of routes it has to the PSTN. These advertisements include fairly dynamic information, such as the remaining capacity in a particular trunk, which is essential for selecting the right gateway.

TGREP is identical in syntax and overall operation to TRIP. However, it differs in the route processing rules followed by the TGREP receiver, allowing for a route processing function called "consolidation". Consolidation combines multiple routes for the same route destination with different attributes to a single route to prevent loss of routing information. TGREP also defines a set of new attributes, usable by TGREP or TRIP. Finally, TGREP only utilizes a subset of overall TRIP capabilities. Specifically, certain attributes are not utilized (as described below), and the TGREP entities (the sender and receiver) operate in an asymmetric relationship, whereas TRIP allows symmetric and asymmetric.

As a general rule, because of a lot of similarities between TRIP and TGREP, frequent reference will be made to the terminologies and formats defined in TRIP [2]. It is suggested that the reader be familiar with the concepts of TRIP like attributes, flags, route types, address families, etc.

2. Terminology and Definitions

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 [1].

Some other useful definitions are:

Circuit: A circuit is a discrete (specific) path between two or more points along which signals can be carried. In this context, a circuit is a physical path, consisting of one or more wires and possibly intermediate switching points.

Trunk: In a network, a trunk is a communication path connecting two switching systems used in the establishment of an end-to-end connection. In selected applications, it may have both its terminations in the same switching system.

TrunkGroup: A trunkgroup is a set of trunks, traffic engineered as a unit, for the establishment of connections within or between switching systems in which all of the paths are interchangeable except where subgrouped.

Carrier: A carrier is a company offering telephone and data communications between points (end-users and/or exchanges).

3. TGREP: Overview of Operation

TGREP is a route registration protocol for telephony destinations on a gateway. These telephony destinations could be prefixes, trunkgroups, or carriers. The TGREP sender resides on the GW and gathers all the information from the GW to relay to the TRIP Location Server. A TGREP Receiver is defined, which receives this information and optionally performs operations like consolidation and aggregation (see Section 7.3), and hands over the reachability information to a TRIP Location Server. The routing proxy also uses the information to select the gateway for incoming calls.

The TGREP sender establishes a session to the TGREP receiver using a procedure similar to session establishment in TRIP. After the session establishment, the TGREP sender sends the reachability information in the UPDATE messages. The UPDATE messages have the same format as in TRIP. However, certain TRIP attributes are not relevant in TGREP; a TGREP receiver MAY ignore them if they are present in a TGREP message. The following TRIP attributes do not apply to TGREP:

1. AdvertisementPath
2. RoutedPath
3. AtomicAggregate
4. LocalPreference
5. MultiExitDisc
6. ITADTopology
7. ConvertedRoute

In addition, TGREP defines the following new attributes in this document that can be carried in a TGREP UPDATE message.

- **TotalCircuitCapacity:** This attribute identifies the total number of PSTN circuits that are available on a route to complete calls.
- **AvailableCircuits:** This attribute identifies the number of PSTN circuits that are currently available on a route to complete calls.

- **CallSuccess:** This attribute represents information about the number of normally terminated calls out of a total number of attempted calls.
- **Prefix (E164):** This attribute is used to represent the list of E164 prefixes to which the respective route can complete calls.
- **Prefix (Decimal Routing Number):** This attribute is used to represent the list of Decimal prefixes to which the respective route can complete calls.
- **Prefix (Hexadecimal Routing Number):** This attribute is used to represent the list of Hexadecimal prefixes to which the respective route can complete calls.
- **TrunkGroup:** This attribute enables providers to route calls to destinations through preferred trunks.
- **Carrier:** This attribute enables providers to route calls to destinations through preferred carriers.

In the rest of the document, we specify attributes and address families used in TGREP. The new attributes and address families introduced are also applicable for general usage in TRIP except where noted (AvailableCircuits attribute, for example).

4. TGREP Attributes

Due to its usage within a service provider network, TGREP makes use of a subset of the attributes defined for TRIP, in addition to defining several new ones. In particular, the following attributes from TRIP are applicable to TGREP:

1. WithdrawnRoutes
2. ReachableRoutes
3. NexthopServer
4. Prefix
5. Communities

TGREP also defines several new attributes, described in this section. These are TotalCircuitCapacity, AvailableCircuits, CallSuccess, TrunkGroup, and Carrier. As mentioned above, these new attributes are usable by TRIP unless noted otherwise.

A TGREP UPDATE supports the following attributes:

1. TotalCircuitCapacity
2. AvailableCircuits
3. CallSuccess
4. Prefix (E.164, Pentadecimal routing number, Decimal routing number)
5. TrunkGroup
6. Carrier

4.1. TotalCircuitCapacity Attribute

Mandatory: False.

Required Flags: Not well-known.

Potential Flags: None.

TRIP Type Code: 13.

The TotalCircuitCapacity attribute identifies the total number of PSTN circuits that are available on a route to complete calls. It is used in conjunction with the AvailableCircuits attribute in gateway selection by the LS. The total number of calls sent to the specified route on the gateway cannot exceed this total circuit capacity under steady state conditions.

The TotalCircuitCapacity attribute is used to reflect the administratively provisioned capacity as opposed to the availability at a given point in time as provided by the AvailableCircuits attribute. Because of its relatively static nature, this attribute MAY be propagated beyond the LS that receives it.

TotalCircuitCapacity represents the total number of possible calls at any instant. This is not expected to change frequently. This can change, for instance, when certain telephony trunks on the gateway are taken out of service for maintenance purposes.

4.1.1. TotalCircuitCapacity Syntax

The TotalCircuitCapacity attribute is a 4-octet unsigned integer. It represents the total number of circuits available for terminating calls through this advertised route. This attribute represents a potentially achievable upper bound on the number of calls that can be terminated on this route in total.

4.1.2. Route Origination and TotalCircuitCapacity

Routes MAY be originated containing the TotalCircuitCapacity attribute.

4.1.3. Route Selection and TotalCircuitCapacity

The TotalCircuitCapacity attribute MAY be used for route selection. Since one of its primary applications is load balancing, an LS will wish to choose a potentially different route (amongst a set of routes for the same destination), on a call-by-call basis. This can be modeled as re-running the decision process on the arrival of each call. The aggregation and dissemination rules for routes with this attribute ensure that re-running this selection process never results in propagation of a new route to other peers.

4.1.4. Aggregation and TotalCircuitCapacity

An LS MAY aggregate routes to the same prefix that contains a TotalCircuitCapacity attribute. It SHOULD add the values of the individual routes to determine the value for the aggregated route in the same ITAD.

4.1.5. Route Dissemination and TotalCircuitCapacity

Since this attribute reflects the static capacity of the gateway's circuit resources, it is not expected to change frequently. Hence, the LS receiving this attribute MAY disseminate it to other peers, both internal and external to the ITAD.

4.2. AvailableCircuits Attribute

Mandatory: False.
Required Flags: Not well-known.
Potential Flags: None.
TRIP Type Code: 14.

The AvailableCircuits attribute identifies the number of PSTN circuits that are currently available on a route to complete calls. The number of additional calls sent to that gateway for that route cannot exceed the circuit capacity. If it does, the signaling protocol will likely generate errors, and calls will be rejected.

The AvailableCircuits attribute is used ONLY between a gateway and its peer LS responsible for managing that gateway. If it is received in a route, it is not propagated.

4.2.1. AvailableCircuits Syntax

The AvailableCircuits attribute is a 4-octet unsigned integer. It represents the number of circuits remaining for terminating calls to this route.

4.2.2. Route Origination and AvailableCircuits

Routes MAY be originated containing the AvailableCircuits attribute. Since this attribute is highly dynamic, changing with every call, updates MAY be sent as it changes. However, it is RECOMMENDED that measures be taken to help reduce the messaging load from route origination. It is further RECOMMENDED that a sufficiently large window of time be used to provide a useful aggregated statistic.

4.2.3. Route Selection and AvailableCircuits

The AvailableCircuits attribute MAY be used for route selection. Since one of its primary applications is load balancing, an LS will wish to choose a potentially different route (amongst a set of routes for the same prefix) on a call-by-call basis. This can be modeled as re-running the decision process on the arrival of each call. The aggregation and dissemination rules for routes with this attribute ensure that re-running this selection process never results in propagation of a new route to other peers.

4.2.4. Aggregation and AvailableCircuits

Not applicable.

4.2.5. Route Dissemination and AvailableCircuits

Routes MUST NOT be disseminated with the AvailableCircuits attribute. The attribute is meant to reflect capacity at the originating gateway only. Its highly dynamic nature makes it inappropriate to disseminate in most cases.

4.3. CallSuccess Attribute

Mandatory: False.

Required Flags: Not well-known.

Potential Flags: None.

TRIP Type Code: 15.

The CallSuccess attribute is an attribute used ONLY between a gateway and its peer LS responsible for managing that gateway. If it is received in a route, it is not propagated.

The CallSuccess attribute provides information about the number of normally terminated calls out of a total number of attempted calls.

CallSuccess is to be determined by the gateway based on the Disconnect cause code at call termination. For example, a call that reaches the Alerting stage but does not get connected due to the

unavailability of the called party, or the called party being busy, is conventionally considered a successful call. On the other hand, a call that gets disconnected because of a circuit or resource being unavailable is conventionally considered a failed call. The exact mapping of disconnect causes to CallSuccess is at the discretion of the gateway reporting the attribute.

The CallSuccess attribute is used by the LS to keep track of failures in reaching certain telephony destinations through a gateway(s) and to use that information in the gateway selection process to enhance the probability of successful call termination.

This information can be used by the LS to consider alternative gateways to terminate calls to those destinations with a better likelihood of success.

4.3.1. CallSuccess Syntax

The CallSuccess attribute is composed of two component fields -- each represented as a 4-octet unsigned integer. The first component field represents the total number of calls terminated successfully for the advertised destination on a given address family over a given window of time. The second component field represents the total number of attempted calls for the advertised destination within the same window of time.

When the value for the total number of attempted calls wraps around, the counter value for the number of successful calls is reset to keep it aligned with the other component over a given window of time. The TGREP receiver using this information should obtain this information frequently enough to prevent loss of significance.

4.3.2. Route Origination and CallSuccess

Routes MAY be originated containing the CallSuccess attribute. This attribute is expected to get statistically significant with passage of time as more calls are attempted. It is RECOMMENDED that sufficiently large windows be used to provide a useful aggregated statistic.

4.3.3. Route Selection and CallSuccess

The CallSuccess attribute MAY be used for route selection. This attribute represents a measure of success of terminating calls to the advertised destination(s). This information MAY be used to select from amongst a set of alternative routes to increase the probability of successful call termination.

4.3.4. Aggregation and CallSuccess

Not applicable.

4.3.5. Route Dissemination and CallSuccess

Routes MUST NOT be disseminated with the CallSuccess attribute. Its potential to change dynamically does not make it suitable to disseminate.

4.4. Prefix Attributes

Mandatory: False.

Required Flags: Not well-known.

Potential Flags: None.

TRIP Type Codes: 16 for E164 Prefix, 17 for Pentadecimal Routing Number Prefix, and 18 for Decimal Routing Number Prefix.

The Prefix attribute is used to represent the list of prefixes to which the respective route can complete calls. This attribute is intended to be used with the Carrier or TrunkGroup address family (discussed in Section 5).

Though it is possible that all prefix ranges may be reachable through a given carrier, administrative issues could make certain ranges preferable to others.

4.4.1. Prefix Attribute Syntax

The Prefix attribute could be E.164, Decimal, or Pentadecimal (refer to TRIP [2]), each of them having its own type code. The Prefix attribute is encoded as a sequence of prefix values in the attribute Value field. The prefixes are listed sequentially with no padding as shown in Figure 2. Each prefix includes a 2-octet Length field that represents the length of the Address field in octets. The order of prefixes in the attribute is not significant.

The presence of the Prefix Attribute with the Length field of the attribute as 0 signifies that the TGREP GW can terminate ALL prefixes of that prefix type (E.164, Decimal, or Pentadecimal) for the given Reachable route(s). This is not equivalent to excluding the Prefix attribute in the TGREP update.




Figure 2: Prefix Format

4.4.2. Route Origination and Prefix

Routes MAY be originated containing the Prefix attribute.

4.4.3. Route Selection and Prefix

The Prefix attribute MAY be used for route selection.

4.4.4. Aggregation and Prefix

Routes with differing Prefix attributes MUST NOT be aggregated.
Routes with the same value in the Prefix attribute MAY be aggregated
and the same Prefix attribute attached to the aggregated object.

4.4.5. Route Dissemination and Prefix

The LS receiving this attribute should disseminate to other peers,
both internal and external to the ITAD.

4.5. TrunkGroup Attribute

Mandatory: False.
Required Flags: Not well-known.
Potential Flags: None.
TRIP Type Code: 19.

The TrunkGroup attribute is used to represent the list of trunkgroups
on the gateway used to complete calls. It enables providers to route
calls to destinations through preferred trunks. This attribute is
relatively static.

4.5.1. TrunkGroup Syntax

The TrunkGroup attribute is a variable-length attribute that is
composed of a sequence of trunkgroup identifiers. It indicates that
the gateway can complete the call through any trunkgroup identifier
indicated in the sequence.

Each trunkgroup identifier is encoded as a Length-Value field (shown
in Figure 3 below). The Length field is a 1-octet unsigned numeric

value. The Value field is a variable-length field consisting of two subfields, a trunkgroup label followed by a trunk context, the two subfields separated by the delimiter ";" (semicolon). Both the trunkgroup label and the trunk context subfields are of variable length. The Length field represents the total size of the Value field including the delimiter.

The permissible character set for the trunk group label and the trunkgroup context subfields and the associated ABNF [8] is per rules outlined in [4].

The presence of the TrunkGroup attribute with the Length field of the attribute as 0 signifies that the TGREP GW can terminate ALL trunkgroups for the given Reachable route(s).

```

< 1 octet > < Length1 octets > < 1 octet > < Length2 octets >

+-----+-----+-----+-----+-----+-----+-----+-----+
| Length1 | TrunkGroup 1 | Length2 | TrunkGroup 2 | ...
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 3: TrunkGroup Syntax

4.5.2. Route Origination and TrunkGroup

Routes MAY be originated containing the TrunkGroup attribute.

4.5.3. Route Selection and TrunkGroup

The TrunkGroup attribute MAY be used for route selection. Certain trunkgroups MAY be preferred over others based on provider policy.

4.5.4. Aggregation and TrunkGroup

Routes with differing TrunkGroup attributes MUST NOT be aggregated. Routes with the same value in the TrunkGroup attribute MAY be aggregated and the same TrunkGroup attribute attached to the aggregated object.

4.5.5. Route Dissemination and TrunkGroup

This attribute is not expected to change frequently. Hence, the LS receiving this attribute MAY disseminate it to other peers, internal to ITAD. Routes SHOULD not be disseminated external to the ITAD, with TrunkGroup attribute.

4.6. Carrier Attribute

Mandatory: False.
 Required Flags: Not well-known.
 Potential Flags: None.
 TRIP Type Code: 20.

The Carrier attribute is used to represent the list of carriers that the gateway uses to complete calls. It enables providers to route calls to destinations through preferred carriers. This attribute is relatively static.

4.6.1. Carrier Syntax

The Carrier attribute is a variable-length attribute that is composed of a sequence of carrier identifiers. It indicates that the route can complete the call to any of the carriers represented in the sequence of carrier identifiers [13].

Each carrier identifier is encoded as a Length-Value field (shown in Figure 4 below). The Length field is a 1-octet unsigned numeric value. The Value field is a variable-length field.

The permissible character set for the Value field and the associated ABNF [8] is per rules outlined in [5]. Specifically, it carries "global-cic" or "local-cic" [5]. In case of "local-cic", the "phonedigit-hex" part and the "cic-context" part would be separated by the delimiter ";". Hence, absence or presence of the delimiter can be used to determine if the value is a "global-cic" or a "local-cic". The Length field represents the total size of the Value field including the delimiter.

The presence of the Carrier attribute with the Length field of the attribute as 0 signifies that the TGREP GW can terminate ALL Carriers for the given Reachable route(s).

```

< 1 octet > < Length1 octets > < 1 octet > < Length2 octets >

+-----+-----+-----+-----+-----+-----+-----+-----+
| Length1 | Carrier 1 | Length2 | Carrier 2 | ... |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 4: Carrier Syntax

4.6.2. Route Origination and Carrier

Routes MAY be originated containing the Carrier attribute.

4.6.3. Route Selection and Carrier

The Carrier attribute MAY be used for route selection. Certain carriers MAY be preferred over others based on provider policy.

4.6.4. Aggregation and Carrier

Routes with differing Carrier attributes MUST NOT be aggregated. Routes with the same value in the Carrier attribute MAY be aggregated and the same Carrier attribute attached to the aggregated object.

4.6.5. Route Dissemination and Carrier

This attribute is not expected to change frequently. Hence, the LS receiving this attribute MAY disseminate it to other peers, both internal and external to the ITAD.

5. TrunkGroup and Carrier Address Families

As described in TRIP [2], the Address Family field gives the type of address for the route. Two new address families and their codes are defined in this section:

Code	Address Family
4	TrunkGroup
5	Carrier

When a set of GWs is to be managed at the granularity of carriers or trunkgroups, then it may be more appropriate for a GW to advertise routes using the Carrier address family or TrunkGroup address family, respectively. Also, in a TGREP association between the gateway and the LS responsible for managing that gateway, there are some attributes that more naturally fit in as advertised properties of trunkgroups or carriers rather than that of advertised prefixes, for example, the AvailableCircuit information on a particular trunkgroup or a particular carrier. To express this relationship, the existing TRIP address families are inadequate. We need separate address families that can associate certain properties like AvailableCircuits information to trunkgroups or carriers.

The primary benefits of this model are as follows:

- It allows a service provider to route calls based strictly on the trunkgroups or carriers.
- It facilitates more accurate reporting of attributes of a dynamic nature like AvailableCircuits by providing the ability to manage resources at the granularity of a trunkgroup or a carrier.

- It enables scalability as gateways can get really large with the ability to provision hundreds or a few thousand circuits, and this can increase the potential for traffic that reports dynamic resource information between the gateway and the LS. The model introduced can potentially alleviate this UPDATE traffic, hence increasing efficiency and providing a scalable gateway registration model.

5.1. Address Family Syntax

Consider the generic TRIP route format from TRIP [2] shown below.




Figure 5: Generic TRIP Route Format

The Address Family field will be used to represent the type of the address associated for this family, which is based on the TrunkGroup or Carrier. The codes for the new address families have been allocated by IANA.

The code for the TrunkGroup address family is 4, and the code for the Carrier address family is 5.

The Application Protocol field is the same as the one for the Decimal, Pentadecimal, and E.164 address families defined in TRIP [2]. The Length field represents the total size of the Address field in bytes.

The value in the Address field represents either the TrunkGroup or Carrier address family, and the syntax is as follows:

For the TrunkGroup address family, the Address field represents a TrunkGroup value that is defined as specified in Section 4.5.1, "TrunkGroup Syntax".

For the Carrier address family, the Address field represents a Carrier value. This is the same as the Value field specified in an earlier Section 4.6.1, "Carrier Syntax".

The above mentioned address families are not hierarchical, but flat. If a gateway supports any of these address families, it should include that address family as one of the Route types supported in the OPEN message capability negotiation phase.

The following attributes are currently defined to be used with all the address families including the TrunkGroup and Carrier address families.

- AvailableCircuits attribute
- TotalCircuitCapacity attribute
- CallSuccess attribute

It is recommended that the above three attributes be used by the gateway with the TrunkGroup or Carrier address family, if possible. This will potentially offer a more efficient resource reporting framework, and a scalable model for gateway provisioning.

However, when the gateway is not using the TrunkGroup or Carrier address family, it MAY use the above attributes with the Decimal, Pentadecimal, and E.164 address families.

The Prefix attribute cannot be used when the address family is E164 numbers, Pentadecimal routing numbers, or Decimal routing numbers.

The Carrier attribute cannot be used if the address family type is Carrier.

The TrunkGroup attribute cannot be used if the address family type is TrunkGroup.

6. Gateway Operation

A gateway uses TGREP to advertise its reachability to its domain's Location Server(s) (LS, which are closely coupled with proxies). The gateway operates in TRIP Send Only mode since it is only interested in advertising its reachability, but is not interested in learning about the reachability of other gateways and other domains. Also, the gateway will not create its own call routing database. In this section, we describe the operation of TGREP on a gateway.

6.1. Session Establishment

When opening a peering session with a TGREP receiver, a TGREP gateway follows exactly the same procedures as any other TRIP entity. The TGREP gateway sends an OPEN message that includes a Send Receive Capability in the Optional Parameters. The Send Receive Capability

is set by the gateway to Send Only. The OPEN message also contains the address families supported by the gateway. The remainder of the peer session establishment is identical to TRIP.

6.2. UPDATE Messages

Once the peer session has been established, the gateway sends UPDATE messages to the TRIP LS with the gateway's entire reachability. The gateway also sends any attributes associated with the routes.

TGREP processing of the UPDATE message at the gateway is identical to UPDATE processing in TRIP [2]. A TGREP sender MUST support all mandatory TRIP attributes.

6.3. KEEPALIVE Messages

KEEPAIVE messages are periodically exchanged over the peering session between the TGREP gateway and the TRIP LS as specified in Section 4.4 of TRIP [2].

6.4. Error Handling and NOTIFICATION Messages

The same procedures used with TRIP are used with TGREP for error handling and generating NOTIFICATION messages. The only difference is that a TGREP gateway will never generate a NOTIFICATION message in response to an UPDATE message, irrespective of the contents of the UPDATE message. Any UPDATE message is silently discarded.

6.5. TGREP Finite State Machine

When the TGREP finite state machine is in the Established state and an UPDATE message is received, the UPDATE message is silently discarded and the TGREP gateway remains in the Established state. Other than that, the TRIP finite state machine and the TGREP finite state machine are identical.

6.6. Call Routing Databases

A TGREP gateway may maintain simultaneous sessions with more than one TRIP LS. A TGREP gateway maintains one call routing database per peer TRIP LS. These databases are equivalent to TRIP's Adj-TRIBS-Out, and hence we will call them Adj-TRIBS-GW-Out. An Adj-TRIB-GW-Out contains the gateway's reachability information advertised to its peer TRIP LS. How an Adj-TRIB-GW-Out database gets populated is outside the scope of this document (possibly by manual configuration).

The TGREP gateway does not have databases equivalent to TRIP's Adj-TRIBs-In and Loc-TRIB, because the TGREP gateway does not learn routes from its peer TRIP LSs, and hence it does not run call route selection.

6.7. Multiple Address Families

As mentioned above, TGREP supports various address families in order to convey the reachability of telephony destinations. A TGREP session MUST NOT send UPDATES of more than one of the following categories (a) Prefix address families (E164, Pentadecimal, and Decimal), (b) TrunkGroup address family, or (c) Carrier address family for a given established session. TGREP should specify its choice address family through the route-type capability in the OPEN message. And route-type specification in the OPEN message violating the above rule should be rejected with a NOTIFICATION message.

6.8. Route Selection and Aggregation

TRIP's route selection and aggregation operations MUST NOT be implemented by TGREP gateways.

7. LS/Proxy Behavior

As mentioned earlier, TGREP can be considered as a protocol complimentary to TRIP in providing reachability information, which can then be further fed into the Location Server. The architecture of an LS/Proxy system is as follows: There exists a TRIP LS application that functions as a speaker in the I-TRIP/E-TRIP network as documented in TRIP [2]. This component is termed as "Egress LS" for the purposes of this discussion. Then, there is a signaling server fronting a set of gateways. In conjunction with this signaling server is also a second component operating in receive mode, which peers with one or more gateways, each of them using TGREP to advertise routing information. This component on the receiving end of one or more TGREP sessions is termed as the "Ingress LS" or "TGREP receiver" for the purposes of this discussion. Also, the entity (typically, a gateway) advertising the routes on the TGREP session is termed as the "TGREP sender". The TGREP receiver receiving the TRIP messages takes the resulting routing information from each gateway, and "exports" it to another process we define below, that performs consolidation and aggregation, in that order. These operations would take as input the collective set of routes from all the gateways. Subsequently, the resulting TRIB is passed as input into the LS-Egress process as shown below, that can then disseminate these via TRIP. The interface between the TGREP receiver (aka Ingress LS) peering with the GW(s) and the TRIP LS (Egress LS) is entirely a local matter.

The nature of the consolidation and aggregation operations and the accompanying motivation are described in the subsections below. The order in which the operations are listed represents an implicit logical sequence in which they are applied. The architecture for an LS/Proxy entity is shown in Figure 6 below.

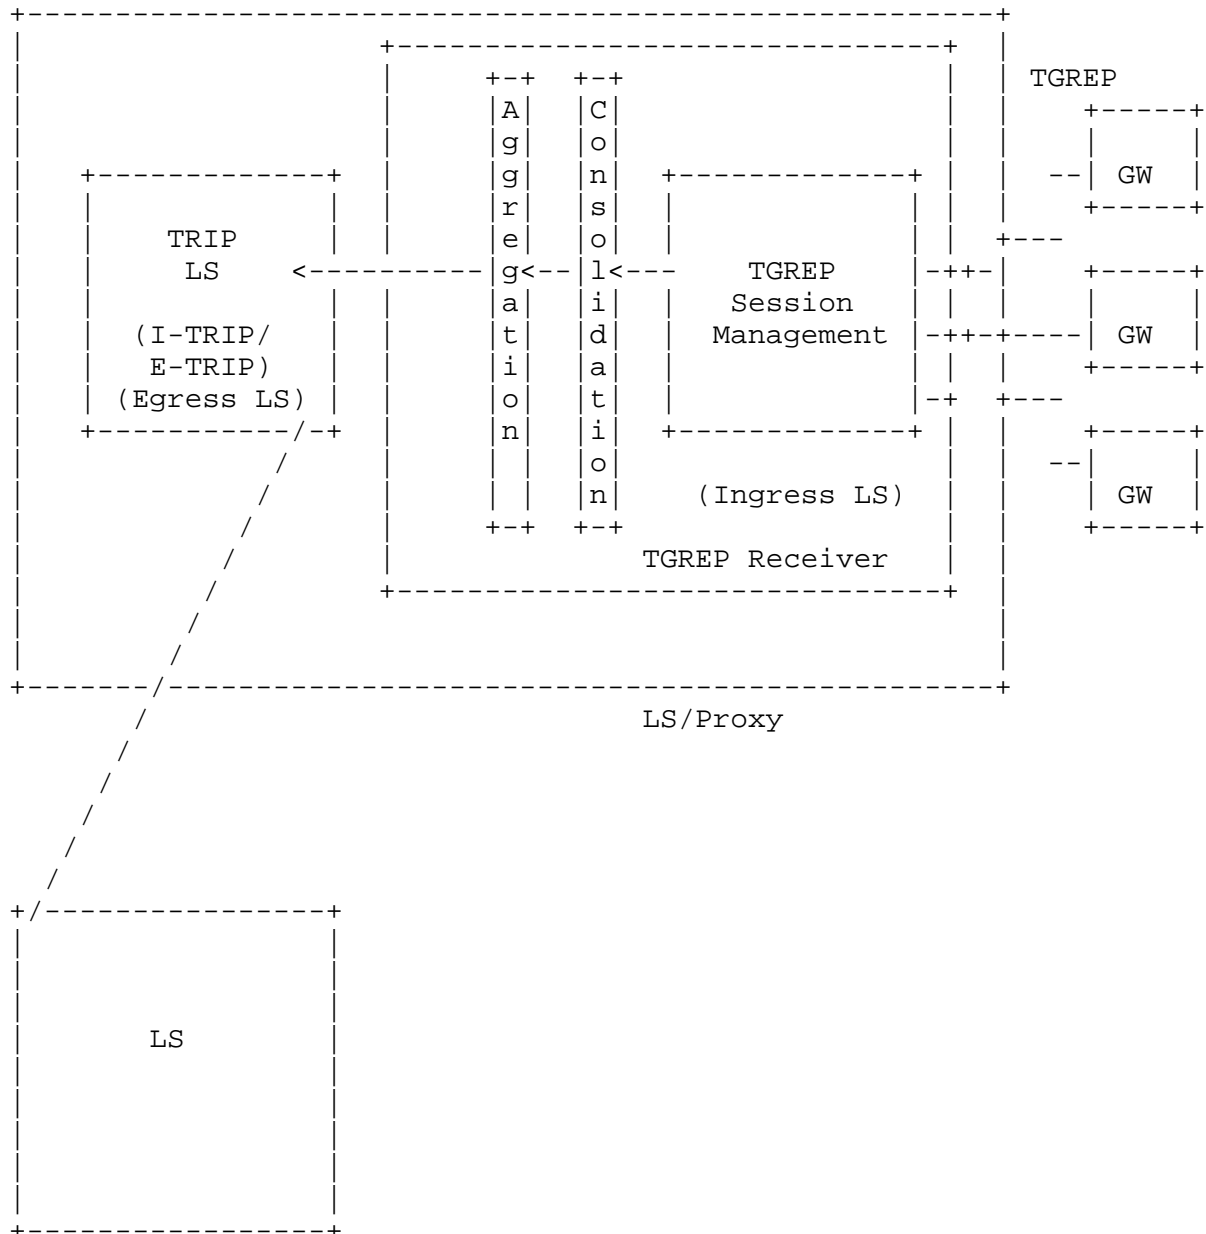


Figure 6: LS Architecture for TRIP-GW

7.1. Route Consolidation

The TGREP receiver (Ingress LS) may receive routing information from one or more gateways. It is possible that multiple routes are available for the same destination. These different alternative routes may be received from the same gateway or from multiple gateways. It is RECOMMENDED that the set of gateway routes for each destination be consolidated, before presenting a candidate route, to the Egress LS entity. The motivation for this operation should be to define a route that can maximally represent the collective routing capabilities of the set of gateways, managed by this TGREP receiver. Let us take an example scenario in order to bring out the motivation for this operation. Let us say, the TGREP receiver maintains peering sessions with gateways A and B.

- Gateway A advertises a route for destination "SIP 408" on the E.164 address family with the Carrier attribute value C1.
- Gateway B advertises a route for destination "SIP 408" on the E.164 address family with Carrier attribute value C2.

The TGREP receiver that receives these routes can consolidate these constituent routes into a single route for destination "SIP 408" with its Carrier attribute being a union of the Carrier attribute values of the individual routes, namely, "C1 C2". This operation is referred to as consolidation. In the above example, it is possible that a route to the destination "SIP 408" through one or more carriers may have been lost if the individual routes were not consolidated.

Another example is to consolidate the Prefix attribute from multiple Carrier or TrunkGroup updates received from different gateways for the same destination. Let us say, there are Carrier address family (AF) updates from two gateways for Carrier destination X, and the prefix attribute values are {408, 650} from one update and {919, 973} from the other. The prefix values from these two updates can be consolidated into a single Carrier AF route advertisement with prefix value {408, 650, 919, 973}.

In general, there is a potential for loss of gateway routing information when TGREP routes from a set of gateways are not consolidated when a candidate route is presented to the TRIP LS. The specifics of applying the consolidation operation to different attributes and routes from different address families is left to the individual TGREP receiver implementations.

7.2. Aggregation

The set of gateway routes, which are in a consolidated form or otherwise, may be aggregated before importing it to the LS instance that is responsible for I-TRIP/E-TRIP processing (Egress LS). This operation follows the standard aggregation procedures described in TRIP [2], while adhering to the aggregation rules for each route attribute.

7.3. Consolidation versus Aggregation

To highlight the difference between the two operations discussed above, "consolidation" combines multiple routes for the same route destination, whereas "aggregation" combines routes for different route destinations that qualify as candidates to be summarized resulting in route information reduction.

To take an example, if there are multiple gateways offering routes to an E.164 destination "408" but with possibly different attributes (e.g., Carrier), the LS/Proxy can combine these to form one route for "408" but representing the attribute information collectively. This process is consolidation.

If, for example, the LS/Proxy receives routes for 4080, 4081, 4082, ... 4089 from amongst a set of gateways, it could aggregate these different candidate routes to have a summarized route destination "408" with each of the attributes computed using the aggregation procedures defined in TRIP.

8. Security Considerations

The security considerations for TGREP are identical to that identified in TRIP [2] and are just restated here for the purposes of clarity.

The security mechanism for the peering session between TGREP GW and a TRIP LS, in an IP network, is IPsec [3]. IPsec uses two protocols to provide traffic security: Authentication Header (AH) [6] and Encapsulating Security Payload (ESP) [7].

The AH header affords data origin authentication, connectionless integrity, and optional anti-replay protection of messages passed between the peer LSs. The ESP header provides origin authentication, connectionless integrity, anti-replay protection, and confidentiality of messages.

Implementations of the protocol defined in this document employing the ESP header SHALL comply with Section 3.1.1 of [10], which defines a minimum set of algorithms for integrity checking and encryption. Similarly, implementations employing the AH header SHALL comply with Section 3.2 of [10], which defines a minimum set of algorithms for integrity checking.

Implementations SHOULD use the Internet Key Exchange Protocol (IKEv2) [9] to permit more robust keying options. Implementations employing IKEv2 SHOULD support 3DES-CBC for confidentiality and HMAC-SHA1 for integrity.

A Security Association (SA) [3] is a simplex "connection" that affords security services to the traffic carried by it. Security services are afforded to an SA by the use of AH or ESP, but not both. Two types of SAs are defined: transport mode and tunnel mode. A transport mode SA is a security association between two hosts, and is appropriate for protecting the TRIP session between two peer LSs.

9. IANA Considerations

Both TRIP [2] and TGREP share the same IANA registry for Capabilities, Attributes, Address Families, and Application Protocols. IANA has added the following attribute codes and address family codes to the TRIP [2] registries.

9.1. Attribute Codes

The Attribute Type Codes assigned for the new attributes defined in this document are listed below:

Code	Attribute	Reference
----	-----	-----
13	TotalCircuitCapacity	[RFC5140]
14	AvailableCircuits	[RFC5140]
15	CallSuccess	[RFC5140]
16	E.164 Prefix	[RFC5140]
17	Pentadecimal Routing Number Prefix	[RFC5140]
18	Decimal Routing Number Prefix	[RFC5140]
19	TrunkGroup	[RFC5140]
20	Carrier	[RFC5140]

9.2. Address Family Codes

The following subsections show the codes that have been assigned for the two new address families introduced in this document.

9.2.1. TrunkGroup Address Family

Code	Address Family	Reference
----	-----	-----
4	TrunkGroup	[RFC5140]

9.2.2. Carrier Address Family

Code	Address Family	Reference
----	-----	-----
5	Carrier	[RFC5140]

10. Acknowledgments

We wish to thank Vijay Gurbani, Li Li, Kevin McDermott, David Oran, Bob Penfield, Jon Peterson, Anirudh Sahoo, and James Yu for their insightful comments and suggestions.

11. References

11.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Rosenberg, J., Salama, H., and M. Squire, "Telephony Routing over IP (TRIP)", RFC 3219, January 2002.
- [3] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [4] Gurbani, V. and C. Jennings, "Representing Trunk Groups in tel/sip Uniform Resource Identifiers (URIs)", RFC 4904, June 2007.
- [5] Yu, J., "Number Portability Parameters for the "tel" URI", RFC 4694, October 2006.
- [6] Kent, S., "IP Authentication Header", RFC 4302, December 2005.
- [7] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.
- [8] Crocker, D., Ed., and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [9] Kaufman, C., Ed., "Internet Key Exchange (IKEv2) Protocol", RFC 4306, December 2005.

- [10] Manral, V., "Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)", RFC 4835, April 2007.

11.2. Informative References

- [11] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [12] Rosenberg, J. and H. Schulzrinne, "A Framework for Telephony Routing over IP", RFC 2871, June 2000.
- [13] ITU-T List of ITU Carrier Codes (published periodically in the ITU-T Operational Bulletin).
- [14] Rosenberg, J., "Requirements for Gateway Registration", Work in Progress, November 2001.

Authors' Addresses

Manjunath S. Bangalore
Cisco
Mail Stop SJC-14/2/1
3625 Cisco Way
San Jose, CA 95134
Phone: +1-408-525-7555
EMail: manjax@cisco.com

Rajneesh Kumar
Cisco
Mail Stop SJC-14/4/2
3625 Cisco Way
San Jose, CA 95134
Phone: +1-408-527-6148
EMail: rajneesh@cisco.com

Jonathan Rosenberg
Cisco
Edison, NJ 08837
EMail: jdrosen@cisco.com

Hussein F. Salama
Citex Software
Giza, Egypt
EMail: hsalama@citexsoftware.com

Dhaval Niranjana Shah
Moowee Inc.
4920 El Camino Real
Los Altos, CA 94022
Phone: +1-408-307-7455
EMail: dhaval@moowee.tv

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

