

Inter-Domain Policy Routing Protocol Specification: Version 1

Status of this Memo

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

Abstract

We present the set of protocols and procedures that constitute Inter-Domain Policy Routing (IDPR). IDPR includes the virtual gateway protocol, the flooding protocol, the route server query protocol, the route generation procedure, the path control protocol, and the data message forwarding procedure.

Contributors

The following people have contributed to the protocols and procedures described in this document: Helen Bowns, Lee Breslau, Ken Carlberg, Isidro Castineyra, Deborah Estrin, Tony Li, Mike Little, Katia Obraczka, Sam Resheff, Martha Steenstrup, Gene Tsudik, and Robert Woodburn.

Table of Contents

1. Introduction.	3
1.1. Domain Elements	3
1.2. Policy.	5
1.3. IDPR Functions.	5
1.3.1. IDPR Entities	6
1.4. Policy Semantics.	7
1.4.1. Source Policies	7
1.4.2. Transit Policies.	8
1.5. IDPR Message Encapsulation.	9
1.5.1. IDPR Data Message Format.	11
1.6. Security.	12
1.7. Timestamps and Clock Synchronization.	13
1.8. Network Management.	14
1.8.1. Policy Gateway Configuration.	17
1.8.2. Route Server Configuration.	18

2. Control Message Transport Protocol.18
2.1. Message Transmission.20
2.2. Message Reception22
2.3. Message Validation.23
2.4. CMTP Message Formats.24
3. Virtual Gateway Protocol.27
3.1. Message Scope28
3.1.1. Pair-PG Messages.28
3.1.2. Intra-VG Messages29
3.1.3. Inter-VG Messages29
3.1.4. VG Representatives.31
3.2. Up/Down Protocol.31
3.3. Implementation.33
3.4. Policy Gateway Connectivity35
3.4.1. Within a Virtual Gateway.35
3.4.2. Between Virtual Gateways.37
3.4.3. Communication Complexity.40
3.5. VGP Message Formats41
3.5.1. UP/DOWN41
3.5.2. PG CONNECT.42
3.5.3. PG POLICY43
3.5.4. VG CONNECT.44
3.5.5. VG POLICY45
3.5.6. Negative Acknowledgements46
4. Routing Information Distribution.47
4.1. AD Representatives.48
4.2. Flooding Protocol48
4.2.1. Message Generation.50
4.2.2. Sequence Numbers.52
4.2.3. Message Acceptance.52
4.2.4. Message Incorporation54
4.2.5. Routing Information Database.56
4.3. Routing Information Message Formats57
4.3.1. CONFIGURATION57
4.3.2. DYNAMIC62
4.3.3. Negative Acknowledgements63
5. Route Server Query Protocol64
5.1. Message Exchange.64
5.2. Remote Route Server Communication65
5.3. Routing Information66
5.4. Routes.67
5.5. Route Server Message Formats.67
5.5.1. ROUTING INFORMATION REQUEST67
5.5.2. ROUTE REQUEST68
5.5.3. ROUTE RESPONSE.71
5.5.4. Negative Acknowledgements72
6. Route Generation.73
6.1. Searching74

6.1.1. Implementation.75
6.2. Route Directionality.78
6.3. Route Database.79
6.3.1. Cache Maintenance80
7. Path Control Protocol and Data Message Forwarding Procedure80
7.1. An Example of Path Setup.81
7.2. Path Identifiers.84
7.3. Path Control Messages85
7.4. Setting Up and Tearing Down a Path.87
7.4.1. Validating Path Identifiers89
7.4.2. Path Consistency with Configured Transit Policies89
7.4.3. Path Consistency with Virtual Gateway Reachability.91
7.4.4. Obtaining Resources92
7.4.5. Target Response93
7.4.6. Originator Response93
7.4.7. Path Life94
7.5. Path Failure and Recovery95
7.5.1. Handling Implicit Path Failures96
7.5.2. Local Path Repair97
7.5.3. Repairing a Path.98
7.6. Path Control Message Formats.	100
7.6.1. SETUP	101
7.6.2. ACCEPT.	103
7.6.3. REFUSE.	103
7.6.4. TEARDOWN.	104
7.6.5. ERROR	105
7.6.6. REPAIR.	106
7.6.7. Negative Acknowledgements	106
8. Security Considerations	106
9. Authors's Address	107
References	107

1. Introduction

In this document, we specify the protocols and procedures that compose Inter-Domain Policy Routing (IDPR). The objective of IDPR is to construct and maintain routes between source and destination administrative domains, that provide user traffic with the services requested within the constraints stipulated for the domains transited. IDPR supports link state routing information distribution and route generation in conjunction with source specified message forwarding. Refer to [5] for a detailed justification of our approach to inter-domain policy routing.

1.1. Domain Elements

The IDPR architecture has been designed to accommodate an internetwork with tens of thousands of administrative domains

collectively containing hundreds of thousands of local networks. Inter-domain policy routes are constructed using information about the services offered by, and the connectivity between, administrative domains. The intra-domain details - gateways, networks, and links traversed - of an inter-domain policy route are the responsibility of intra-domain routing and are thus outside the scope of IDPR.

An "administrative domain" (AD) is a collection of contiguous hosts, gateways, networks, and links managed by a single administrative authority. The domain administrator defines service restrictions for transit traffic and service requirements for locally-generated traffic, and selects the addressing schemes and routing procedures that apply within the domain. Within the Internet, each domain has a unique numeric identifier assigned by the Internet Assigned Numbers Authority (IANA).

"Virtual gateways" (VGs) are the only IDPR-recognized connecting points between adjacent domains. Each virtual gateway is a collection of directly-connected "policy gateways" (see below) in two adjoining domains, whose existence has been sanctioned by the administrators of both domains. The domain administrators may agree to establish more than one virtual gateway between the two domains. For each such virtual gateway, the two administrators together assign a local numeric identifier, unique within the set of virtual gateways connecting the two domains. To produce a virtual gateway identifier unique within its domain, a domain administrator concatenates the mutually assigned local virtual gateway identifier together with the adjacent domain's identifier.

Policy gateways (PGs) are the physical gateways within a virtual gateway. Each policy gateway enforces service restrictions on IDPR transit traffic, as stipulated by the domain administrator, and forwards the traffic accordingly. Within a domain, two policy gateways are "neighbors" if they are in different virtual gateways. A single policy gateway may belong to multiple virtual gateways. Within a virtual gateway, two policy gateways are "peers" if they are in the same domain and are "adjacent" if they are in different domains. Adjacent policy gateways are "directly connected" if the only Internet-addressable entities attached to the connecting medium are policy gateways in the virtual gateways. Note that this definition implies that not only point-to-point links but also networks may serve as direct connections between adjacent policy gateways. The domain administrator assigns to each of its policy gateways a numeric identifier, unique within that domain.

A "domain component" is a subset of a domain's entities such that all entities within the subset are mutually reachable via intra-domain routes, but no entities outside the subset are reachable via intra-

domain routes from entities within the subset. Normally, a domain consists of a single component, namely itself; however, when partitioned, a domain consists of multiple components. Each domain component has an identifier, unique within the Internet, composed of the domain identifier together with the identifier of the lowest-numbered operational policy gateway within the component. All operational policy gateways within a domain component can discover mutual reachability through intra-domain routing information. Hence, all such policy gateways can consistently determine, without explicit negotiation, which of them has the lowest number.

1.2. Policy

With IDPR, each domain administrator sets "transit policies" that dictate how and by whom the resources in its domain should be used. Transit policies are usually public, and they specify offered services comprising:

- Access restrictions: e.g., applied to traffic to or from certain domains or classes of users.
- Quality: e.g., delay, throughput, or error characteristics.
- Monetary cost: e.g., charge per byte, message, or unit time.

Each domain administrator also sets "source policies" for traffic originating in its domain. Source policies are usually private, and they specify requested services comprising:

- Access restrictions: e.g., domains to favor or avoid in routes.
- Quality: e.g., acceptable delay, throughput, and reliability.
- Monetary cost: e.g., acceptable session cost.

1.3. IDPR Functions

IDPR comprises the following functions:

- Collecting and distributing routing information including domain transit policies and inter-domain connectivity.
- Generating and selecting policy routes based on the routing information distributed and on the source policies configured or requested.
- Setting up paths across the Internet using the policy routes generated.

- Forwarding messages across and between domains along the established paths.
- Maintaining databases of routing information, inter-domain policy routes, forwarding information, and configuration information.

1.3.1. IDPR Entities

Several different entities are responsible for performing the IDPR functions.

Policy gateways, the only IDPR-recognized connecting points between adjacent domains, collect and distribute routing information, participate in path setup, forward data messages along established paths, and maintain forwarding information databases.

"Path agents", resident within policy gateways and within "route servers" (see below), act on behalf of hosts to select policy routes, to set up and manage paths, and to maintain forwarding information databases. Any Internet host can reap the benefits of IDPR, as long as there exists a path agent configured to act on its behalf and a means by which the host's messages can reach the path agent. Specifically, a path agent in one domain may be configured to act on behalf of hosts in another domain. In this case, the path agent's domain is an IDPR "proxy" for the hosts' domain.

Route servers maintain both the routing information database and the route database, and they generate policy routes using the routing information collected and the source policies requested by the path agents. A route server may reside within a policy gateway, or it may exist as an autonomous entity. Separating the route server functions from the policy gateways frees the policy gateways from both the memory intensive task of database (routing information and route) maintenance and the computationally intensive task of route generation. Route servers, like policy gateways, each have a unique numeric identifier within their domain, assigned by the domain administrator.

Given the size of the current Internet, each policy gateway can perform the route server functions, in addition to its message forwarding functions, with little or no degradation in message forwarding performance. Aggregating the routing functions into policy gateways simplifies implementation; one need only install IDPR protocols in policy gateways. Moreover, it simplifies communication between routing functions, as all functions reside within each policy gateway. As the Internet grows, the memory and processing required to perform the route server functions may become a burden for the policy gateways. When this happens, each domain administrator should

separate the route server functions from the policy gateways in its domain.

"Mapping servers" maintain the database of mappings that resolve Internet names and addresses to domain identifiers. Each host is contained within a domain and is associated with a proxy domain which may be identical with the host's domain. The mapping server function will be integrated into the existing DNS name service (see [6]) and will provide mappings between a host and its local and proxy domains.

"Configuration servers" maintain the databases of configured information that apply to IDPR entities within their domains. Configuration information for a given domain includes transit policies (i.e., service offerings and restrictions), source policies (i.e., service requirements), and mappings between local IDPR entities and their names and addresses. The configuration server function will be integrated into a domain's existing network management system (see [7]-[8]).

1.4. Policy Semantics

The source and transit policies supported by IDPR are intended to accommodate a wide range of services available throughout the Internet. We describe the semantics of these policies, concentrating on the access restriction aspects. To express these policies in this document, we have chosen to use a syntactic variant of Clark's policy term notation [1]. However, we provide a more succinct syntax (see [7]) for actually configuring source and transit policies.

1.4.1. Source Policies

Each source policy takes the form of a collection of sets as follows:

Applicable Sources and Destinations:

$\{((H(1,1),s(1,1)),\dots,(H(1,f1),s(1,f1))),\dots,((H(n,1),s(n,1)),\dots,(H(n,fn),s(n,fn)))\}$: The set of groups of source/destination traffic flows to which the source policy applies. Each traffic flow group $((H(i,1),s(i,1)),\dots,(H(i,fi),s(i,fi)))$ contains a set of source hosts and corresponding destination hosts. Here, $H(i,j)$ represents a host, and $s(i,j)$, an element of $\{\text{SOURCE}, \text{DESTINATION}\}$, represents an indicator of whether $H(i,j)$ is to be considered as a source or as a destination.

Domain Preferences: $\{(AD(1),x(1)),\dots,(AD(m),x(m))\}$: The set of transit domains that the traffic flows should favor, avoid, or exclude. Here, $AD(i)$ represents a domain, and $x(i)$, an element of $\{\text{FAVOR}, \text{AVOID}, \text{EXCLUDE}\}$, represents an indicator of whether routes including $AD(i)$ are to be favored, avoided if possible, or

unconditionally excluded.

UCI: The source user class for the traffic flows listed.

RequestedServices: The set of requested services not related to access restrictions, i.e., service quality and monetary cost.

When selecting a route for a traffic flow from a source host $H(i,j)$ to a destination host $H(i,k)$, where $1 \leq i \leq n$ and $1 \leq j, k \leq f_i$, the path agent (see section 1.3.1) must honor the source policy such that:

- For each domain, $AD(p)$, contained in the route, $AD(p)$ is not equal to any $AD(k)$, such that $1 \leq k \leq m$ and $x(k) = EXCLUDE$.
- The route provides the services listed in the set Requested Services.

1.4.2. Transit Policies

Each transit policy takes the form of a collection of sets as follows:

Source/Destination Access Restrictions:

$\{((H(1,1),AD(1,1),s(1,1)),\dots,(H(1,f_1),AD(1,f_1),s(1,f_1))),\dots,((H(n,1),AD(n,1),s(n,1)),\dots,(H(n,f_n),AD(n,f_n),s(n,f_n)))\}$: The set of groups of source and destination hosts and domains to which the transit policy applies. Each domain group $((H(i,1),AD(i,1),s(i,1)),\dots,(H(i,f_i),AD(i,f_i),s(i,f_i)))$ contains a set of source and destination hosts and domains such that this transit domain will carry traffic from each source listed to each destination listed. Here, $H(i,j)$ represents a set of hosts, $AD(i,j)$ represents a domain containing $H(i,j)$, and $s(i,j)$, a subset of $\{SOURCE, DESTINATION\}$, represents an indicator of whether $(H(i,j),AD(i,j))$ is to be considered as a set of sources, destinations, or both.

Temporal Access Restrictions: The set of time intervals during which the transit policy applies.

User Class Access Restrictions: The set of user classes to which the transit policy applies.

Offered Services: The set of offered services not related to access restrictions, i.e., service quality and monetary cost.

Virtual Gateway Access Restrictions:

$\{((VG(1,1),e(1,1)),\dots,(VG(1,g_1),e(1,g_1))),\dots,((VG(m,1),e(m,1)),\dots,(VG(m,g_m),e(m,g_m)))\}$ gateways to which the transit policy applies. Each virtual gateway group $((VG(i,1),e(i,1)),\dots,(VG(i,g_i),e(i,g_i)))$ contains a set of domain entry and exit points such that each entry virtual gateway can reach (barring an intra-domain routing failure) each exit virtual gateway via an intra-domain route supporting the transit policy. Here, $VG(i,j)$ represents a virtual gateway, and $e(i,j)$, a subset of $\{ENTRY, EXIT\}$, represents an indicator of whether $VG(i,j)$ is to be considered as a domain entry point, exit point, or both.

The domain advertising such a transit policy will carry traffic from any host in the set $H(i,j)$ in $AD(i,j)$ to any host in the set $H(i,k)$ in $AD(i,k)$, where $1 \leq i \leq n$ and $1 \leq j, k \leq f_i$, provided that:

- SOURCE is an element of $s(i,j)$.
- DESTINATION is an element of $s(i,k)$.
- Traffic from $H(i,j)$ enters the domain during one of the intervals in the set Temporal Access Restrictions.
- Traffic from $H(i,j)$ carries one of the user class identifiers in the set User Class Access Restrictions.
- Traffic from $H(i,j)$ enters via any $VG(u,v)$ such that ENTRY is an element of $e(u,v)$, where $1 \leq u \leq m$ and $1 \leq v \leq g_u$.
- Traffic to $H(i,k)$ leaves via any $VG(u,w)$ such that EXIT is an element of $e(u,w)$, where $1 \leq w \leq g_u$.

1.5. IDPR Message Encapsulation

There are two kinds of IDPR messages:

- "Data messages" containing user data generated by hosts.
- "Control messages" containing IDPR protocol-related control information generated by policy gateways and route servers.

Within an internetwork, only policy gateways and route servers are able to generate, recognize, and process IDPR messages. The existence of IDPR is invisible to all other gateways and hosts, including mapping servers and configuration servers. Mapping servers and configuration servers perform necessary but ancillary functions

for IDPR, and thus they are not required to handle IDPR messages.

An IDPR entity places IDPR-specific information in each IDPR control message it originates; this information is significant only to recipient IDPR entities. Using "encapsulation" across each domain, an IDPR message tunnels from source to destination across an internetwork through domains that may employ disparate intra-domain addressing schemes and routing procedures.

As an alternative to encapsulation, we had considered embedding IDPR in IP, as a set of IP options. However, this approach has the following disadvantages:

- Only domains that support IP would be able to participate in IDPR; domains that do not support IP would be excluded.
- Each gateway, policy or other, in a participating domain would at least have to recognize the IDPR option, even if it did not execute the IDPR protocols. However, most commercial routers are not optimized for IP options processing, and so IDPR message handling might require significant processing at each gateway.
- For some IDPR protocols, in particular path control, the size restrictions on IP options would preclude inclusion of all of the necessary protocol-related information.

For these reasons, we decided against the IP option approach and in favor of encapsulation.

An IDPR message travels from source to destination between consecutive policy gateways. Each policy gateway encapsulates the IDPR message with information, for example an IP header, that will enable the message to reach the next policy gateway. Note that the encapsulating header and the IDPR-specific information may increase the message size beyond the MTU of the given domain. However, message fragmentation and reassembly is the responsibility of the protocol, for example IP, that encapsulates IDPR messages for transport between successive policy gateways; it is not currently the responsibility of IDPR itself.

A policy gateway, when forwarding an IDPR message to a peer or a neighbor policy gateway, encapsulates the message in accordance with the addressing scheme and routing procedure of the given domain and indicates in the protocol field of the encapsulating header that the message is indeed an IDPR message. Intermediate gateways between the two policy gateways forward the IDPR message as they would any other message, using the information in the encapsulating header. Only the recipient policy gateway interprets the protocol field, strips off

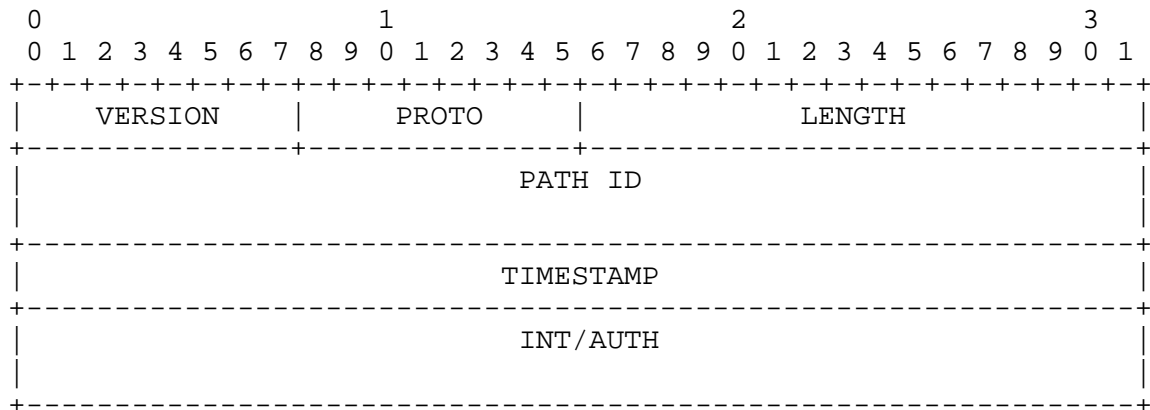
the encapsulating header, and processes the IDPR message.

A policy gateway, when forwarding an IDPR message to a directly-connected adjacent policy gateway, encapsulates the message in accordance with the addressing scheme of the entities within the virtual gateway and indicates in the protocol field of the encapsulating header that the message is indeed an IDPR message. The recipient policy gateway strips off the encapsulating header and processes the IDPR message. We recommend that the recipient policy gateway perform the following validation check of the encapsulating header, prior to stripping it off. Specifically, the recipient policy gateway should verify that the source address and the destination address in the encapsulating header match the adjacent policy gateway's address and its own address, respectively. Moreover, the recipient policy gateway should verify that the message arrived on the interface designated for the direct connection to the adjacent policy gateway. These checks help to ensure that IDPR traffic that crosses domain boundaries does so only over direct connections between adjacent policy gateways.

Policy gateways forward IDPR data messages according to a forwarding information database which maps "path identifiers", carried in the data messages, into next policy gateways. Policy gateways forward IDPR control messages according to next policy gateways selected by the particular IDPR control protocols associated with the messages. Distinguishing IDPR data messages and IDPR control messages at the encapsulating protocol level, instead of at the IDPR protocol level, eliminates an extra level of dispatching and hence makes IDPR message forwarding more efficient. When encapsulated within IP messages, IDPR data messages and IDPR control messages carry the IP protocol numbers 35 and 38, respectively.

1.5.1. IDPR Data Message Format

The path agents at a source domain determine which data messages generated by local hosts are to be handled by IDPR. To each data message selected for IDPR handling, a source path agent prepends the following header:



VERSION (8 bits) Version number for IDPR data messages, currently equal to 1.

PROTO (8 bits) Numeric identifier for the protocol with which to process the contents of the IDPR data message. Only the path agent at the destination interprets and acts upon the contents of the PROTO field.

LENGTH (16 bits) Length of the entire IDPR data message in bytes.

PATH ID (64 bits) Path identifier assigned by the source's path agent and consisting of the numeric identifier for the path agent's domain (16 bits), the numeric identifier for the path agent's policy gateway (16 bits), and the path agent's local path identifier (32 bits) (see section 7.2).

TIMESTAMP (32 bits) Number of seconds elapsed since 1 January 1970 0:00 GMT.

INT/AUTH (variable) Computed integrity/authentication value, dependent on the type of integrity/authentication requested during path setup.

We describe the IDPR control message header in section 2.4.

1.6. Security

IDPR contains mechanisms for verifying message integrity and source authenticity and for protecting against certain types of denial of service attacks. It is particularly important to keep IDPR control messages intact, because they carry control information critical to the construction and use of viable policy routes between domains.

All IDPR messages carry a single piece of information, referred to as

the "integrity/authentication value", which may be used not only to detect message corruption but also to verify the authenticity of the message source. In the Internet, the IANA will sanction the set of valid algorithms which may be used to compute the integrity/authentication values. This set may include algorithms that perform only message integrity checks such as n-bit cyclic redundancy checksums (CRCs), as well as algorithms that perform both message integrity and source authentication checks such as signed hash functions of message contents.

Each domain administrator is free to select any integrity/authentication algorithm, from the set specified by the IANA, for computing the integrity/authentication values contained in its domain's messages. However, we recommend that IDPR entities in each domain be capable of executing all of the valid algorithms so that an IDPR control message originating at an entity in one domain can be properly checked by an entity in another domain.

Each IDPR control message must carry a non-null integrity/authentication value. We recommend that control message integrity/authentication be based on a digital signature algorithm applied to a one-way hash function, such as RSA applied to MD5 [17], which simultaneously verifies message integrity and source authenticity. The digital signature may be based on either public-key or private-key cryptography. Our approach to digital signature use in IDPR is based on the privacy-enhanced Internet electronic mail service [13]-[15], already available in the Internet.

We do not require that IDPR data messages carry a non-null integrity/authentication value. In fact, we recommend that a higher layer (end-to-end) procedure, and not IDPR, assume responsibility for checking the integrity and authenticity of data messages, because of the amount of computation involved.

1.7. Timestamps and Clock Synchronization

Each IDPR message carries a timestamp (expressed in seconds elapsed since 1 January 1970 0:00 GMT, following the UNIX precedent) supplied by the source IDPR entity, which serves to indicate the age of the message. IDPR entities use the absolute value of the timestamp to confirm that a message is current and use the relative difference between timestamps to determine which message contains the more recent information.

All IDPR entities must possess internal clocks that are synchronized to some degree, in order for the absolute value of a message timestamp to be meaningful. The synchronization granularity required by IDPR is on the order of minutes and can be achieved manually.

Thus, a clock synchronization protocol operating among all IDPR entities in all domains, while useful, is not necessary.

An IDPR entity can determine whether to accept or reject a message based on the discrepancy between the message's timestamp and the entity's own internal clock time. Any IDPR message whose timestamp lies outside of the acceptable range may contain stale or corrupted information or may have been issued by a source whose internal clock has lost synchronization with the message recipient's internal clock. Timestamp checks are required for control messages because of the consequences of propagating and acting upon incorrect control information. However, timestamp checks are discretionary for data messages but may be invoked during problem diagnosis, for example, when checking for suspected message replays.

We note that none of the IDPR protocols contain explicit provisions for dealing with an exhausted timestamp space. As timestamp space exhaustion will not occur until well into the next century, we expect timestamp space viability to outlast the IDPR protocols.

1.8. Network Management

In this document, we do not describe how to configure and manage IDPR. However, in this section, we do provide a list of the types of IDPR configuration information required. Also, in later sections describing the IDPR protocols, we briefly note the types of exceptional events that must be logged for network management. Complete descriptions of IDPR entity configuration and IDPR managed objects appear in [7] and [8] respectively.

To participate in inter-domain policy routing, policy gateways and route servers within a domain each require configuration information. Some of the configuration information is specifically defined within the given domain, while some of the configuration information is universally defined throughout an internetwork. A domain administrator determines domain-specific information, and in the Internet, the IANA determines globally significant information.

To produce valid domain configurations, the domain administrators must receive the following global information from the IANA:

- For each integrity/authentication type, the numeric identifier, syntax, and semantics. Available integrity and authentication types include but are not limited to:
 - o public-key based signatures;
 - o private-key based signatures;

- o cyclic redundancy checksums;
 - o no integrity/authentication.
- For each user class, the numeric identifier, syntax, and semantics. Available user classes include but are not limited to:
 - o federal (and if necessary, agency-specific such as NSF, DOD, DOE, etc.);
 - o research;
 - o commercial;
 - o support.
- For each offered service that may be advertised in transit policies, the numeric identifier, syntax, and semantics. Available offered services include but are not limited to:
 - o average message delay;
 - o message delay variation;
 - o average bandwidth available;
 - o available bandwidth variation;
 - o maximum transfer unit (MTU);
 - o charge per byte;
 - o charge per message;
 - o charge per unit time.
- For each access restriction that may be advertised in transit policies, the numeric identifier, syntax, and semantics. Available access restrictions include but are not limited to:
 - o Source and destination domains and host sets.
 - o User classes.
 - o Entry and exit virtual gateways.
 - o Time of day.

- For each requested service that may appear within a path setup message, the numeric identifier, syntax, and semantics. Available requested services include but are not limited to:
 - o maximum path life in minutes, messages, or bytes;
 - o integrity/authentication algorithms to be used on data messages sent over the path;
 - o upper bound on path delay;
 - o minimum delay path;
 - o upper bound on path delay variation;
 - o minimum delay variation path;
 - o lower bound on path bandwidth;
 - o maximum bandwidth path;
 - o upper bound on monetary cost;
 - o minimum monetary cost path.

In an internetwork-wide implementation of IDPR, the set of global configuration parameters and their syntax and semantics must be consistent across all participating domains. The IANA, responsible for establishing the full set of global configuration parameters in the Internet, relies on the cooperation of the administrators of all participating domains to ensure that the global parameters are consistent with the desired transit policies and user service requirements of each domain. Moreover, as the syntax and semantics of the global parameters affects the syntax and semantics of the corresponding IDPR software, the IANA must carefully define each global parameter so that it is unlikely to require future modification.

The IANA provides configured global information to configuration servers in all domains participating in IDPR. Each domain administrator uses the configured global information maintained by its configuration servers to develop configurations for each IDPR entity within its domain. Each configuration server retains a copy of the configuration for each local IDPR entity and also distributes the configuration to that entity using, for example, SNMP.

1.8.1. Policy Gateway Configuration

Each policy gateway must contain sufficient configuration information to perform its IDPR functions, which subsume those of the path agent. These include: validating IDPR control messages; generating and distributing virtual gateway connectivity and routing information messages to peer, neighbor, and adjacent policy gateways; distributing routing information messages to route servers in its domain; resolving destination addresses; requesting policy routes from route servers; selecting policy routes and initiating path setup; ensuring consistency of a path with its domain's transit policies; establishing path forwarding information; and forwarding IDPR data messages along existing paths. The necessary configuration information includes the following:

- For each integrity/authentication type, the numeric identifier, syntax, and semantics.
- For each policy gateway and route server in the given domain, the numeric identifier and set of addresses or names.
- For each virtual gateway connected to the given domain, the numeric identifier, the numeric identifiers for the constituent peer policy gateways, and the numeric identifier for the adjacent domain.
- For each virtual gateway of which the given policy gateway is a member, the numeric identifiers and set of addresses for the constituent adjacent policy gateways.
- For each policy gateway directly-connected and adjacent to the given policy gateway, the local connecting interface.
- For each local route server to which the given policy gateway distributes routing information, the numeric identifier.
- For each source policy applicable to hosts within the given domain, the syntax and semantics.
- For each transit policy applicable to the domain, the numeric identifier, syntax, and semantics.
- For each requested service that may appear within a path setup message, the numeric identifier, syntax, and semantics.
- For each source user class, the numeric identifier, syntax, and semantics.

1.8.2. Route Server Configuration

Each route server must contain sufficient configuration information to perform its IDPR functions, which subsume those of the path agent. These include: validating IDPR control messages; deciphering and storing the contents of routing information messages; exchanging routing information with other route servers and policy gateways; generating policy routes that respect transit policy restrictions and source service requirements; distributing policy routes to path agents in policy gateways; resolving destination addresses; selecting policy routes and initiating path setup; establishing path forwarding information; and forwarding IDPR data messages along existing paths. The necessary configuration information includes the following:

- For each integrity/authentication type, the numeric identifier, syntax, and semantics.
- For each policy gateway and route server in the given domain, the numeric identifier and set of addresses or names.
- For each source policy applicable to hosts within the given domain, the syntax and semantics.
- For access restriction that may be advertised in transit policies, the numeric identifier, syntax, and semantics.
- For each offered service that may be advertised in transit policies, the numeric identifier, syntax, and semantics.
- For each requested service that may appear within a path setup message, the numeric identifier, syntax, and semantics.
- For each source user class, the numeric identifier, syntax, and semantics.

2. Control Message Transport Protocol

IDPR control messages convey routing-related information that directly affects the policy routes generated and the paths set up across the Internet. Errors in IDPR control messages can have widespread, deleterious effects on inter-domain policy routing, and so the IDPR protocols have been designed to minimize loss and corruption of control messages. For every control message it transmits, each IDPR protocol expects to receive notification as to whether the control message successfully reached the intended IDPR recipient. Moreover, the IDPR recipient of a control message first verifies that the message appears to be well-formed, before acting on its contents.

All IDPR protocols use the Control Message Transport Protocol (CMTTP), a connectionless, transaction-based transport layer protocol, for communication with intended recipients of control messages. CMTTP retransmits unacknowledged control messages and applies integrity and authenticity checks to received control messages.

There are three types of CMTTP messages:

DATAGRAM:

Contains IDPR control messages.

ACK: Positive acknowledgement in response to a DATAGRAM message.

NAK: Negative acknowledgement in response to a DATAGRAM message.

Each CMTTP message contains several pieces of information supplied by the sender that allow the recipient to test the integrity and authenticity of the message. The set of integrity and authenticity checks performed after CMTTP message reception are collectively referred to as "validation checks" and are described in section 2.3.

When we first designed the IDPR protocols, CMTTP as a distinct protocol did not exist. Instead, CMTTP-equivalent functionality was embedded in each IDPR protocol. To provide a cleaner implementation, we later decided to provide a single transport protocol that could be used by all IDPR protocols. We originally considered using an existing transport protocol, but rejected this approach for the following reasons:

- The existing reliable transport protocols do not provide all of the validation checks, in particular the timestamp and authenticity checks, required by the IDPR protocols. Hence, if we were to use one of these protocols, we would still have to provide a separate protocol on top of the transport protocol to force retransmission of IDPR messages that failed to pass the required validation checks.
- Many of the existing reliable transport protocols are window-based and hence can result in increased message delay and resource use when, as is the case with IDPR, multiple independent messages use the same transport connection. A single message experiencing transmission problems and requiring retransmission can prevent the window from advancing, forcing all subsequent messages to queue behind it. Moreover, many of the window-based protocols do not support selective retransmission of failed messages but instead require retransmission of not only the failed message but also all preceding messages within the window.

For these reasons, we decided against using an existing transport

protocol and in favor of developing CMTTP.

2.1. Message Transmission

At the transmitting entity, when an IDPR protocol is ready to issue a control message, it passes a copy of the message to CMTTP; it also passes a set of parameters to CMTTP for inclusion in the CMTTP header and for proper CMTTP message handling. In turn, CMTTP converts the control message and associated parameters into a DATAGRAM by prepending the appropriate header to the control message. The CMTTP header contains several pieces of information to aid the message recipient in detecting errors (see section 2.4). Each IDPR protocol can specify all of the following CMTTP parameters applicable to its control message:

- IDPR protocol and message type.
- Destination.
- Integrity/authentication scheme.
- Timestamp.
- Maximum number of transmissions allotted.
- Retransmission interval in microseconds.

One of these parameters, the timestamp, can be specified directly by CMTTP as the internal clock time at which the message is transmitted. However, two of the IDPR protocols, namely flooding and path control, themselves require message generation timestamps for proper protocol operation. Thus, instead of requiring CMTTP to pass back a timestamp to an IDPR protocol, we simplify the service interface between CMTTP and the IDPR protocols by allowing an IDPR protocol to specify the timestamp in the first place.

Using the control message and accompanying parameters supplied by the IDPR protocol, CMTTP constructs a DATAGRAM, adding to the header CMTTP-specific parameters. In particular, CMTTP assigns a "transaction identifier" to each DATAGRAM generated, which it uses to associate acknowledgements with DATAGRAM messages. Each DATAGRAM recipient includes the received transaction identifier in its returned ACK or NAK, and each DATAGRAM sender uses the transaction identifier to match the received ACK or NAK with the original DATAGRAM.

A single DATAGRAM, for example a routing information message or a path control message, may be handled by CMTTP at many different policy gateways. Within a pair of consecutive IDPR entities, the DATAGRAM

sender expects to receive an acknowledgement from the DATAGRAM recipient. However, only the IDPR entity that actually generated the original CMTP DATAGRAM has control over the transaction identifier, because that entity may supply a digital signature that covers the entire DATAGRAM. The intermediate policy gateways that transmit the DATAGRAM do not change the transaction identifier. Nevertheless, at each DATAGRAM recipient, the transaction identifier must uniquely distinguish the DATAGRAM so that only one acknowledgement from the next DATAGRAM recipient matches the original DATAGRAM. Therefore, the transaction identifier must be globally unique.

The transaction identifier consists of the numeric identifiers for the domain and IDPR entity (policy gateway or route server) issuing the original DATAGRAM, together with a 32-bit local identifier assigned by CMTP operating within that IDPR entity. We recommend implementing the 32-bit local identifier either as a simple counter incremented for each DATAGRAM generated or as a fine granularity clock. The former always guarantees uniqueness of transaction identifiers; the latter guarantees uniqueness of transaction identifiers, provided the clock granularity is finer than the minimum possible interval between DATAGRAM generations and the clock wrapping period is longer than the maximum round-trip delay to and from any internetwork destination.

Before transmitting a DATAGRAM, CMTP computes the length of the entire message, taking into account the prescribed integrity/authentication scheme, and then computes the integrity/authentication value over the whole message. CMTP includes both of these quantities, which are crucial for checking message integrity and authenticity at the recipient, in the DATAGRAM header. After sending a DATAGRAM, CMTP saves a copy and sets an associated retransmission timer, as directed by the IDPR protocol parameters. If the retransmission timer fires and CMTP has received neither an ACK nor a NAK for the DATAGRAM, CMTP then retransmits the DATAGRAM, provided this retransmission does not exceed the transmission allotment. Whenever a DATAGRAM exhausts its transmission allotment, CMTP discards the DATAGRAM, informs the IDPR protocol that the control message transmission was not successful, and logs the event for network management. In this case, the IDPR protocol may either resubmit its control message to CMTP, specifying an alternate destination, or discard the control message altogether.

2.2. Message Reception

At the receiving entity, when CMTTP obtains a DATAGRAM, it takes one of the following actions, depending upon the outcome of the message validation checks:

- The DATAGRAM passes the CMTTP validation checks. CMTTP then delivers the DATAGRAM with enclosed IDPR control message, to the appropriate IDPR protocol, which in turn applies its own integrity checks to the control message before acting on the contents. The recipient IDPR protocol, except in one case, directs CMTTP to generate an ACK and return the ACK to the sender. That exception is the up/down protocol (see section 3.2) which determines reachability of adjacent policy gateways and does not use CMTTP ACK messages to notify the sender of message reception. Instead, the up/down protocol messages themselves carry implicit information about message reception at the adjacent policy gateway. In the cases where the recipient IDPR protocol directs CMTTP to generate an ACK, it may pass control information to CMTTP for inclusion in the ACK, depending on the contents of the original IDPR control message. For example, a route server unable to fill a request for routing information may inform the requesting IDPR entity, through an ACK for the initial request, to place its request elsewhere.
- The DATAGRAM fails at least one of the CMTTP validation checks. CMTTP then generates a NAK, returns the NAK to the sender, and discards the DATAGRAM, regardless of the type of IDPR control message contained in the DATAGRAM. The NAK indicates the nature of the validation failure and serves to help the sender establish communication with the recipient. In particular, the CMTTP NAK provides a mechanism for negotiation of IDPR version and integrity/authentication scheme, two parameters crucial for establishing communication between IDPR entities.

Upon receiving an ACK or a NAK, CMTTP immediately discards the message if at least one of the validation checks fails or if it is unable to locate the associated DATAGRAM. CMTTP logs the latter event for network management. Otherwise, if all of the validation checks pass and if it is able to locate the associated DATAGRAM, CMTTP clears the associated retransmission timer and then takes one of the following actions, depending upon the message type:

- The message is an ACK. CMTTP discards the associated DATAGRAM and delivers the ACK, which may contain IDPR control information, to the appropriate IDPR protocol.
- The message is a NAK. If the associated DATAGRAM has exhausted its transmission allotment, CMTTP discards the DATAGRAM, informs the

appropriate IDPR protocol that the control message transmission was not successful, and logs the event for network management. Otherwise, if the associated DATAGRAM has not yet exhausted its transmission allotment, CMTP first checks its copy of the DATAGRAM against the failure indication contained in the NAK. If its DATAGRAM copy appears to be intact, CMTP retransmits the DATAGRAM and sets the associated retransmission timer. However, if its DATAGRAM copy appears to be corrupted, CMTP discards the DATAGRAM, informs the IDPR protocol that the control message transmission was not successful, and logs the event for network management.

2.3. Message Validation

On every CMTP message received, CMTP performs a set of validation checks to test message integrity and authenticity. The order in which these tests are executed is important. CMTP must first determine if it can parse enough of the message to compute the integrity/authentication value. (Refer to section 2.4 for a description of CMTP message formats.) Then, CMTP must immediately compute the integrity/authentication value before checking other header information. An incorrect integrity/authentication value means that the message is corrupted, and so it is likely that CMTP header information is incorrect. Checking specific header fields before computing the integrity/authentication value not only may waste time and resources, but also may lead to incorrect diagnoses of a validation failure.

The CMTP validation checks are as follows:

- CMTP verifies that it can recognize both the control message version type contained in the header. Failure to recognize either one of these values means that CMTP cannot continue to parse the message.
- CMTP verifies that it can recognize and accept the integrity/authentication type contained in the header; no integrity/authentication is not an acceptable type for CMTP.
- CMTP computes the integrity/authentication value and verifies that it equals the integrity/authentication value contained in the header. For key-based integrity/authentication schemes, CMTP may use the source domain identifier contained in the CMTP header to index the correct key. Failure to index a key means that CMTP cannot compute the integrity/authentication value.
- CMTP computes the message length in bytes and verifies that it equals the length value contained in the header.

- CMTTP verifies that the message timestamp is in the acceptable range. The message should be no more recent than `cmtp_new` (300) seconds ahead of the entity's current internal clock time. In this document, when we present an IDPR system configuration parameter, such as `cmtp_new`, we usually follow it with a recommended value in parentheses. The `cmtp_new` value allows some clock drift between IDPR entities. Moreover, each IDPR protocol has its own limit on the maximum age of its control messages. The message should be no less recent than a prescribed number of seconds behind the recipient entity's current internal clock time. Hence, each IDPR protocol performs its own message timestamp check in addition to that performed by CMTTP.
- CMTTP verifies that it can recognize the IDPR protocol designated for the enclosed control message.


Whenever CMTTP encounters a failure while performing any of these validation checks, it logs the event for network management. If the failure occurs on a DATAGRAM, CMTTP immediately generates a NAK containing the reason for the failure, returns the NAK to the sender, and discards the DATAGRAM message. If the failure occurs on an ACK or a NAK, CMTTP discards the ACK or NAK message.

2.4. CMTTP Message Formats

In designing the format of IDPR control messages, we have attempted to strike a balance between efficiency of link bandwidth usage and efficiency of message processing. In general, we have chosen compact representations for IDPR information in order to minimize the link bandwidth consumed by IDPR-specific information. However, we have also organized IDPR information in order to speed message processing, which does not always result in minimum link bandwidth usage.

To limit link bandwidth usage, we currently use fixed-length identifier fields in IDPR messages; domains, virtual gateways, policy gateways, and route servers are all represented by fixed-length identifiers. To simplify message processing, we currently align fields containing an even number of bytes on even-byte boundaries within a message. In the future, if the Internet adopts the use of super domains, we will offer hierarchical, variable-length identifier fields in an updated version of IDPR.

The header of each CMTTP message contains the following information:

**VERSION**

(8 bits) Version number for IDPR control messages, currently equal to 1.

PRT (4 bits) Numeric identifier for the control message transport protocol, equal to 0 for CMTTP.

MSG (4 bits) Numeric identifier for the CMTTP message type, equal to 0 for a DATAGRAM, 1 for an ACK, and 2 for a NAK.

DPR (4 bits) Numeric identifier for the original DATAGRAM's IDPR protocol type.

DMS (4 bits) Numeric identifier for the original DATAGRAM's IDPR message type.

I/A TYP (8 bits) Numeric identifier for the integrity/authentication scheme used. CMTTP requires the use of an integrity/authentication scheme; this value must not be set equal to 0, indicating no integrity/authentication in use.

SOURCE AD (16 bits) Numeric identifier for the domain containing the IDPR entity that generated the message.

SOURCE ENT (16 bits) Numeric identifier for the IDPR entity that generated the message.

TRANSACTION ID (32 bits) Local transaction identifier assigned by the IDPR entity that generated the original DATAGRAM.

TIMESTAMP (32 bits) Number of seconds elapsed since 1 January 1970 0:00 GMT.

LENGTH (16 bits) Length of the entire IDPR control message, including the CMTTP header, in bytes.

message specific (16 bits) Dependent upon CMTTP message type.

For DATAGRAM and ACK messages:

RESERVED

(16 bits) Reserved for future use and currently set equal to 0.

For NAK messages:

ERR TYP (8 bits) Numeric identifier for the type of CMTTP validation failure encountered. Validation failures include the following types:

1. Unrecognized IDPR control message version number.
2. Unrecognized CMTTP message type.
3. Unrecognized integrity/authentication scheme.
4. Unacceptable integrity/authentication scheme.
5. Unable to locate key using source domain.
6. Incorrect integrity/authentication value.
7. Incorrect message length.
8. Message timestamp out of range.
9. Unrecognized IDPR protocol designated for the enclosed control message.

ERR INFO (8 bits) CMTTP supplies the following additional information for the designated types of validation failures:

Type 1:

Acceptable IDPR control message version number.

Types 3 and 4: Acceptable integrity/authentication type.

DATAGRAM AD

(16 bits) Numeric identifier for the domain containing the IDPR entity that generated the original DATAGRAM. Present only in ACK and NAK messages.

DATAGRAM ENT (16 bits) Numeric identifier for the IDPR entity that generated the original DATAGRAM. Present only in ACK and NAK messages.

INFORM (optional, variable) Information to be interpreted by the IDPR protocol that issued the original DATAGRAM. Present only in ACK messages and dependent on the original DATAGRAM's IDPR protocol type.

INT/AUTH (variable) Computed integrity/authentication value, dependent on the type of integrity/authentication scheme used.

3. Virtual Gateway Protocol

Every policy gateway within a domain participates in gathering information about connectivity within and between virtual gateways of which it is a member and in distributing this information to other virtual gateways in its domain. We refer to these functions collectively as the Virtual Gateway Protocol (VGP).

The information collected through VGP has both local and global significance for IDPR. Virtual gateway connectivity information, distributed to policy gateways within a single domain, aids those policy gateways in selecting routes across and between virtual gateways connecting their domain to adjacent domains. Inter-domain connectivity information, distributed throughout an internetwork in routing information messages, aids route servers in constructing feasible policy routes.

Provided that a domain contains simple virtual gateway and transit policy configurations, one need only implement a small subset of the VGP functions. The connectivity among policy gateways within a virtual gateway and the heterogeneity of transit policies within a

domain determine which VGP functions must be implemented, as we explain toward the end of this section.

3.1. Message Scope

Policy gateways generate VGP messages containing information about perceived changes in virtual gateway connectivity and distribute these messages to other policy gateways within the same domain and within the same virtual gateway. We classify VGP messages into three distinct categories: "pair-PG", "intra-VG", and "inter-VG", depending upon the scope of message distribution.

Policy gateways use CMTTP for reliable transport of VGP messages. The issuing policy gateway must communicate to CMTTP the maximum number of transmissions per VGP message, `vgp_ret`, and the interval between VGP message retransmissions, `vgp_int` microseconds. The recipient policy gateway must determine VGP message acceptability; conditions of acceptability depend on the type of VGP message, as we describe below.

Policy gateways store, act upon, and in the case of inter-VG messages, forward the information contained in acceptable VGP messages. VGP messages that pass the CMTTP validation checks but fail a specific VGP message acceptability check are considered to be unacceptable and are hence discarded by recipient policy gateways. A policy gateway that receives an unacceptable VGP message also logs the event for network management.

3.1.1. Pair-PG Messages

Pair-PG message communication occurs between the two members of a pair of adjacent, peer, or neighbor policy gateways. With IDPR, the only pair-PG messages are those periodically generated by the up/down protocol and used to monitor mutual reachability between policy gateways.

A pair-PG message is "acceptable" if:

- It passes the CMTTP validation checks.
- Its timestamp is less than `vgp_old` (300) seconds behind the recipient's internal clock time.
- Its destination policy gateway identifier coincides with the identifier of the recipient policy gateway.
- Its source policy gateway identifier coincides with the identifier of a policy gateway configured for the recipient's domain or

associated virtual gateway.

3.1.2. Intra-VG Messages

Intra-VG message communication occurs between one policy gateway and all of its peers. Whenever a policy gateway discovers that its connectivity to an adjacent or neighbor policy gateway has changed, it issues an intra-VG message indicating the connectivity change to all of its reachable peers. Whenever a policy gateway detects that a previously unreachable peer is now reachable, it issues, to that peer, intra-VG messages indicating connectivity to adjacent and neighbor policy gateways. If the issuing policy gateway fails to receive an analogous intra-VG message from the newly reachable peer within twice the configured VGP retransmission interval, `vgp_int` microseconds, it actively requests the intra-VG message from that peer. These message exchanges ensure that peers maintain a consistent view of each others' connectivity to adjacent and neighbor policy gateways.

An intra-VG message is "acceptable" if:

- It passes the CMTP validation checks.
- Its timestamp is less than `vgp_old` (300) seconds behind the recipient's internal clock time.
- Its virtual gateway identifier coincides with that of a virtual gateway configured for the recipient's domain.

3.1.3. Inter-VG Messages

Inter-VG message communication occurs between one policy gateway and all of its neighbors. Whenever the lowest-numbered operational policy gateway in a set of mutually reachable peers discovers that its virtual gateway's connectivity to the adjacent domain or to another virtual gateway has changed, it issues an inter-VG message indicating the connectivity change to all of its neighbors. Specifically, the policy gateway distributes an inter-VG message to a "VG representative" policy gateway (see section 3.1.4 below) in each virtual gateway in the domain. Each VG representative in turn propagates the inter-VG message to each of its peers.

Whenever the lowest-numbered operational policy gateway in a set of mutually peers detects that one or more previously unreachable peers are now reachable, it issues, to the lowest-numbered operational policy gateway in all other virtual gateways, requests for inter-VG information indicating connectivity to adjacent domains and to other virtual gateways. The recipient policy gateways return the requested

inter-VG messages to the issuing policy gateway, which in turn distributes the messages to the newly reachable peers. These message exchanges ensure that virtual gateways maintain a consistent view of each others' connectivity, while consuming minimal domain resources in distributing connectivity information.

An inter-VG message contains information about the entire virtual gateway, not just about the issuing policy gateway. Thus, when virtual gateway connectivity changes happen in rapid succession, recipients of the resultant inter-VG messages should be able to determine the most recent message and that message must contain the current virtual gateway connectivity information. To ensure that the connectivity information distributed is consistent and unambiguous, we designate a single policy gateway, namely the lowest-numbered operational peer, for generating and distributing inter-VG messages. It is a simple procedure for a set of mutually reachable peers to determine the lowest-numbered member, as we describe in section 3.2 below.

To understand why a single member of a virtual gateway must issue inter-VG messages, consider the following example. Suppose that two peers in a virtual gateway each detect a different connectivity change and generate separate inter-VG messages. Recipients of these messages may not be able to determine which message is more recent if policy gateway internal clocks are not perfectly synchronized. Moreover, even if the clocks were perfectly synchronized, and hence message recency could be consistently determined, it is possible for each peer to issue its inter-VG message before receiving current information from the other. As a result, neither inter-VG message contains the correct connectivity from the perspective of the virtual gateway. However, these problems are eliminated if all inter-VG messages are generated by a single peer within a virtual gateway, in particular the lowest-numbered operational policy gateway.

An inter-VG message is "acceptable" if:

- It passes the CMTP validation checks.
- Its timestamp is less than vgp_old (300) seconds behind the recipient's internal clock time.
- Its virtual gateway identifier coincides with that of a virtual gateway configured for the recipient's domain.
- Its source policy gateway identifier represents the lowest numbered operational member of the issuing virtual gateway, reachable from the recipient.

Distribution of intra-VG messages among peers often triggers generation and distribution of inter-VG messages among virtual gateways. Usually, the lowest-numbered operational policy gateway in a virtual gateway generates and distributes an inter-VG message immediately after detecting a change in virtual gateway connectivity, through receipt or generation of an intra-VG message. However, if this policy gateway is also waiting for an intra-VG message from a newly reachable peer, it does not immediately generate and distribute the inter-VG message.

Waiting for intra-VG messages enables the lowest-numbered operational policy gateway in a virtual gateway to gather the most recent connectivity information for inclusion in the inter-VG message. However, under unusual circumstances, the policy gateway may fail to receive an intra-VG message from a newly reachable peer, even after actively requesting such a message. To accommodate this case, VGP uses an upper bound of four times the configured retransmission interval, `vgp_int` microseconds, on the amount of time to wait before generating and distributing an inter-VG message, when receipt of an intra-VG message is pending.

3.1.4. VG Representatives

When distributing an inter-VG message, the issuing policy gateway selects as recipients one neighbor, the VG Representative, from each virtual gateway in the domain. To be selected as a VG representative, a policy gateway must be reachable from the issuing policy gateway via intra-domain routing. The issuing policy gateway gives preference to neighbors that are members of more than one virtual gateway. Such a neighbor acts as a VG representative for all virtual gateways of which it is a member and restricts inter-VG message distribution as follows: any policy gateway that is a peer in more than one of the represented virtual gateways receives at most one copy of the inter-VG message. This message distribution strategy minimizes the number of message copies required for disseminating inter-VG information.

3.2. Up/Down Protocol

Directly-connected adjacent policy gateways execute the Up/Down Protocol to determine mutual reachability. Pairs of peer or neighbor policy gateways can determine mutual reachability through information provided by the intra-domain routing procedure or through execution of the up/down protocol. In general, we do not recommend implementing the up/down protocol between each pair of policy gateways in a domain, as it results in $O(n^2)$ (where n is the number of policy gateways within the domain) communications complexity. However, if the intra-domain routing procedure is slow to detect

connectivity changes or is unable to report reachability at the IDPR entity level, the reachability information obtained through the up/down protocol may well be worth the extra communications cost. In the remainder of this section, we describe the up/down protocol from the perspective of adjacent policy gateways, but we note that the identical protocol can be applied to peer and neighbor policy gateways as well.

The up/down protocol determines whether the direct connection between adjacent policy gateways is acceptable for data traffic transport. A direct connection is presumed to be "down" (unacceptable for data traffic transport) until the up/down protocol declares it to be "up" (acceptable for data traffic transport). We say that a virtual gateway is "up" if there exists at least one pair of adjacent policy gateways whose direct connection is acceptable for data traffic transport, and that a virtual gateway is "down" if there exists no such pair of adjacent policy gateways.

When executing the up/down protocol, policy gateways exchange UP/DOWN messages every `ud_per` (1) second. All policy gateways use the same default period of `ud_per` initially and then negotiate a preferred period through exchange of UP/DOWN messages. A policy gateway reports its desired value for `ud_per` within its UP/DOWN messages. It then chooses the larger of its desired value and that of the adjacent policy gateway as the period for exchanging subsequent UP/DOWN messages. Policy gateways also exchange, in UP/DOWN messages, information about the identity of their respective domain components. This information assists the policy gateways in selecting routes across virtual gateways to partitioned domains.

Each UP/DOWN message is transported using CMTTP and hence is covered by the CMTTP validation checks. However, unlike other IDPR control messages, UP/DOWN messages do not require reliable transport. Specifically, the up/down protocol requires only a single transmission per UP/DOWN message and never directs CMTTP to return an ACK. As pair-PG messages, UP/DOWN messages are acceptable under the conditions described in section 3.1.1.

Each policy gateway assesses the state of its direct connection, to the adjacent policy gateway, by counting the number of acceptable UP/DOWN messages received within a set of consecutive periods. A policy gateway communicates its perception of the state of the direct connection through its UP/DOWN messages. Initially, a policy gateway indicates the down state in each of its UP/DOWN messages. Only when the direct connection appears to be up from its perspective does a policy gateway indicate the up state in its UP/DOWN messages.

A policy gateway can begin to transport data traffic over a direct

connection only if both of the following conditions are true:

- The policy gateway receives from the adjacent policy gateway at least j acceptable UP/DOWN messages within the last m consecutive periods. From the recipient policy gateway's perspective, this event up. Hence, the recipient policy gateway indicates the up state in its subsequent UP/DOWN messages.
- The UP/DOWN message most recently received from the adjacent policy gateway indicates the up state, signifying that the adjacent policy gateway considers the direct connection to be up.

A policy gateway must cease to transport data traffic over a direct connection whenever either of the following conditions is true:

- The policy gateway receives from the adjacent policy gateway at most acceptable UP/DOWN messages within the last n consecutive periods.
- The UP/DOWN message most recently received from the adjacent policy gateway indicates the down state, signifying that the adjacent policy gateway considers the direct connection to be down.

From the recipient policy gateway's perspective, either of these events constitutes a state transition of the direct connection from up to down. Hence, the policy gateway indicates the down state in its subsequent UP/DOWN messages.

3.3. Implementation

We recommend implementing the up/down protocol using a sliding window. Each window slot indicates the UP/DOWN message activity during a given period, containing either a "hit" for receipt of an acceptable UP/DOWN message or a "miss" for failure to receive an acceptable UP/DOWN message. In addition to the sliding window, the implementation should include a tally of hits recorded during the current period and a tally of misses recorded over the current window.

When the direct connection moves to the down state, the initial values of the up/down protocol parameters must be set as follows:

- The sliding window size is equal to m .
- Each window slot contains a miss.
- The current period hit tally is equal to 0.

- The current window miss tally is equal to m .

When the direct connection moves to the up state, the initial values of the up/down protocol parameters must be set as follows:

- The sliding window size is equal to n .
- Each window slot contains a hit.
- The current period hit tally is equal to 0.
- The current window miss tally is equal to 0.

At the conclusion of each period, a policy gateway computes the miss tally and determines whether there has been a state transition of the direct connection to the adjacent policy gateway. In the down state, a miss tally of no more than $m - j$ signals a transition to the up state. In the up state, a miss tally of no less than $n - k$ signals a transition to the down state.

Computing the correct miss tally involves several steps. First, the policy gateway prepares to slide the window by one slot so that the oldest slot disappears, making room for the newest slot. However, before sliding the window, the policy gateway checks the contents of the oldest window slot. If this slot contains a miss, the policy gateway decrements the miss tally by 1, as this slot is no longer part of the current window.

After sliding the window, the policy gateway determines the proper contents. If the hit tally for the current period equals 0, the policy gateway records a miss for the newest slot and increments the miss tally by 1. Otherwise, if the hit tally for the current period is greater than 0, the policy gateway records a hit for the newest slot and decrements the hit tally by 1. Moreover, the policy gateway applies any remaining hits to slots containing misses, beginning with the newest and progressing to the oldest such slot. For each such slot containing a miss, the policy gateway records a hit in that slot and decrements both the hit and miss tallies by 1, as the hit cancels out a miss. The policy gateway continues to apply each remaining hit tallied to any slot containing a miss, until either all such hits are exhausted or all such slots are accounted for. Before beginning the next up/down period, the policy gateway resets the hit tally to 0.

Although we expect the hit tally, within any given period, to be no greater than 1, we do anticipate the occasional period in which a policy gateway receives more than one UP/DOWN message from an adjacent policy gateway. The most common reasons for this occurrence are message delay and clock drift. When an UP/DOWN message is

delayed, the receiving policy gateway observes a miss in one period followed by two hits in the next period, one of which cancels the previous miss. However, excess hits remaining in the tally after miss cancellation indicate a problem, such as clock drift. Thus, whenever a policy gateway accumulates excess hits, it logs the event for network management.

When clock drift occurs between two adjacent policy gateways, it causes the period of one policy gateway to grow with respect to the period of the other policy gateway. Let $p(X)$ be the period for PG X, let $p(Y)$ be the period for PG Y, and let g and h be the smallest positive integers such that $g * p(X) = h * p(Y)$. Suppose that $p(Y) > p(X)$ because of clock drift. In this case, PG X observes $g - h$ misses in g consecutive periods, while PG Y observes $g - h$ surplus hits in h consecutive periods. As long as $(g - h)/g < (n - k)/n$ and $(g - h)/g < \text{or} = (m - j)/m$, the clock drift itself will not cause the direct connection to enter or remain in the down state.

3.4. Policy Gateway Connectivity

Policy gateways collect connectivity information through the intra-domain routing procedure and through VGP, and they distribute connectivity changes through VGP in both intra-VG messages to peers and inter-VG messages to neighbors. Locally, this connectivity information assists policy gateways in selecting routes, not only across a virtual gateway to an adjacent domain but also across a domain between two virtual gateways. Moreover, changes in connectivity between domains are distributed, in routing information messages, to route servers throughout an internetwork.

3.4.1. Within a Virtual Gateway

Each policy gateway within a virtual gateway constantly monitors its connectivity to all adjacent and to all peer policy gateways. To determine the state of its direct connection to an adjacent policy gateway, a policy gateway uses reachability information supplied by the up/down protocol. To determine the state of its intra-domain routes to a peer policy gateway, a policy gateway uses reachability information supplied by either the intra-domain routing procedure or the up/down protocol.

A policy gateway generates a PG CONNECT message whenever either of the following conditions is true:

- The policy gateway detects a change, in state or in adjacent domain component, associated with its direct connection to an adjacent policy gateway. In this case, the policy gateway distributes a copy of the message to each peer reachable via

intra-domain routing.

- The policy gateway detects that a previously unreachable peer is now reachable. In this case, the policy gateway distributes a copy of the message to the newly reachable peer.

A PG CONNECT message is an intra-VG message that includes information about each adjacent policy gateway directly connected to the issuing policy gateway. Specifically, the PG CONNECT message contains the adjacent policy gateway's identifier, status (reachable or unreachable), and domain component identifier. If a PG CONNECT message contains a "request", each peer that receives the message responds to the sender with its own PG CONNECT message.

All mutually reachable peers monitor policy gateway connectivity within their virtual gateway, through the up/down protocol, the intra-domain routing procedure, and the exchange of PG CONNECT messages. Within a given virtual gateway, each constituent policy gateway maintains the following information about each configured adjacent policy gateway:

- The identifier for the adjacent policy gateway.
- The status of the adjacent policy gateway: reachable/unreachable, directly connected/not directly connected.
- The local exit interfaces used to reach the adjacent policy gateway, provided it is reachable.
- The identifier for the adjacent policy gateway's domain component.
- The set of peers to which the adjacent policy gateway is directly-connected.

Hence, all mutually reachable peers can detect changes in connectivity across the virtual gateway to adjacent domain components.

When the lowest-numbered operational peer policy gateway within a virtual gateway detects a change in the set of adjacent domain components reachable through direct connections across the given virtual gateway, it generates a VGCONNECT message and distributes a copy to a VG representative in all other virtual gateways connected to its domain. A VG CONNECT message is an inter-VG message that includes information about each peer's connectivity across the given virtual gateway. Specifically, the VG CONNECT message contains, for each peer, its identifier and the identifiers of the domain components reachable through its direct connections to adjacent

policy gateways. Moreover, the VG CONNECT message gives each recipient enough information to determine the state, up or down, of the issuing virtual gateway.

The issuing policy gateway, namely the lowest-numbered operational peer, may have to wait up to four times `vgp_int` microseconds after detecting the connectivity change, before generating and distributing the VGCONNECT message, as described in section 3.1.3. Each recipient VG representative in turn distributes a copy of the VG CONNECT message to each of its peers reachable via intra-domain routing. If a VG CONNECT message contains a "request", then in each recipient virtual gateway, the lowest-numbered operational peer that receives the message responds to the original sender with its own VGCONNECT message.

3.4.2. Between Virtual Gateways

At present, we expect transit policies to be uniform over all intra-domain routes between any pair of policy gateways within a domain. However, when tariffed qualities of service become prevalent offerings for intra-domain routing, we can no longer expect uniformity of transit policies throughout a domain. To monitor the transit policies supported on intra-domain routes between virtual gateways requires both a policy-sensitive intra-domain routing procedure and a VGP exchange of policy information between neighbor policy gateways.

Each policy gateway within a domain constantly monitors its connectivity to all peer and neighbor policy gateways, including the transit policies supported on intra-domain routes to these policy gateways. To determine the state of its intra-domain connection to a peer or neighbor policy gateway, a policy gateway uses reachability information supplied by either the intra-domain routing procedure or the up/down protocol. To determine the transit policies supported on intra-domain routes to a peer or neighbor policy gateway, a policy gateway uses policy-sensitive reachability information supplied by the intra-domain routing procedure. We note that when transit policies are uniform over a domain, reachability and policy-sensitive reachability are equivalent.

Within a virtual gateway, each constituent policy gateway maintains the following information about each configured peer and neighbor policy gateway:

- The identifier for the peer or neighbor policy gateway.
- The identifiers corresponding to the transit policies configured to be supported by intra-domain routes to the peer or neighbor policy

gateway.

- According to each transit policy, the status of the peer or neighbor policy gateway: reachable/unreachable.
- For each transit policy, the local exit interfaces used to reach the peer or neighbor policy gateway, provided it is reachable.
- The identifiers for the adjacent domain components reachable through direct connections from the peer or neighbor policy gateway, obtained through VG CONNECT messages.

Using this information, a policy gateway can detect changes in its connectivity to an adjoining domain component, with respect to a given transit policy and through a given neighbor. Moreover, combining the information obtained for all neighbors within a given virtual gateway, the policy gateway can detect changes in its connectivity, with respect to a given transit policy, to that virtual gateway and to adjoining domain components reachable through that virtual gateway.

All policy gateways mutually reachable via intra-domain routes supporting a configured transit policy need not exchange information about perceived changes in connectivity, with respect to the given transit policy. In this case, each policy gateway can infer another's policy-sensitive reachability to a third, through mutual intra-domain reachability information provided by the intra-domain routing procedure. However, whenever two or more policy gateways are no longer mutually reachable with respect to a given transit policy, these policy gateways can no longer infer each other's reachability to other policy gateways, with respect to that transit policy. In this case, these policy gateways must exchange explicit information about changes in connectivity to other policy gateways, with respect to that transit policy.

A policy gateway generates a PG POLICY message whenever either of the following conditions is true:

- The policy gateway detects a change in its connectivity to another virtual gateway, with respect to a configured transit policy, or to an adjoining domain component reachable through that virtual gateway. In this case, the policy gateway distributes a copy of the message to each peer reachable via intra-domain routing but not currently reachable via any intra-domain routes of the given transit policy.
- The policy gateway detects that a previously unreachable peer is reachable. In this case, the policy gateway distributes a copy of

the message to the newly reachable peer.

A PG POLICY message is an intra-VG message that includes information about each configured transit policy and each virtual gateway configured to be reachable from the issuing policy gateway via intra-domain routes of the given transit policy. Specifically, the PGPOLICY message contains, for each configured transit policy:

- The identifier for the transit policy.
- The identifiers for the virtual gateways associated with the given transit policy and currently reachable, with respect to that transit policy, from the issuing policy gateway.
- The identifiers for the domain components reachable from and adjacent to the members of the given virtual gateways.

If a PG POLICY message contains a "request", each peer that receives the message responds to the original sender with its own PG POLICY message.

In addition to connectivity between itself and its neighbors, each policy gateway also monitors the connectivity, between domain components adjacent to its virtual gateway and domain components adjacent to other virtual gateways, through its domain and with respect to the configured transit policies. For each member of each of its virtual gateways, a policy gateway monitors:

- The set of adjacent domain components currently reachable through direct connections across the given virtual gateway. The policy gateway obtains this information through PG CONNECT messages from reachable peers and through UP/DOWN messages from adjacent policy gateways.
- For each configured transit policy, the set of virtual gateways currently reachable from the given virtual gateway with respect to that transit policy and the set of adjoining domain components currently reachable through direct connections across those virtual gateways. The policy gateway obtains this information through PG POLICY messages from peers, VG CONNECT messages from neighbors, and the intra-domain routing procedure. Using this information, a policy gateway can detect connectivity changes, through its domain and with respect to a given transit policy, between adjoining domain components.

When the lowest-numbered operational policy gateway within a virtual gateway detects a change in the connectivity between a domain component adjacent to its virtual gateway and a domain component

adjacent to another virtual gateway in its domain, with respect to a configured transit policy, it generates a VG POLICY message and distributes a copy to a VG representative in selected virtual gateways connected to its domain. In particular, the lowest-numbered operational policy gateway distributes a VG POLICY message to a VG representative in every other virtual gateway containing a member reachable via intra-domain routing but not currently reachable via any routes of the given transit policy. A VG POLICY message is an inter-VG message that includes information about the connectivity between domain components adjacent to the issuing virtual gateway and domain components adjacent to the other virtual gateways in the domain, with respect to configured transit policies. Specifically, the VG POLICY message contains, for each transit policy:

- The identifier for the transit policy.
- The identifiers for the virtual gateways associated with the given transit policy and currently reachable, with respect to that transit policy, from the issuing virtual gateway.
- The identifiers for the domain components reachable from and adjacent to the members of the given virtual gateways.

The issuing policy gateway, namely the lowest-numbered operational peer, may have to wait up to four times `vgp_int` microseconds after detecting the connectivity change, before generating and distributing the VG POLICY message, as described in section 3.1.3. Each recipient VG representative in turn distributes a copy of the VG POLICY message to each of its peers reachable via intra-domain routing. If a VG POLICY message contains a "request", then in each recipient virtual gateway, the lowest-numbered operational peer that receives the message responds to the original sender with its own VG POLICY message.

3.4.3. Communication Complexity

We offer an example, to provide an estimate of the number of VGP messages exchanged within a domain, AD X, after a detected change in policy gateway connectivity. Suppose that an adjacent domain, AD Y, partitions such that the partition is detectable through the exchange of UP/DOWN messages across a virtual gateway connecting AD X and AD Y. Let *V* be the number of virtual gateways in AD X. Suppose each virtual gateway contains *P* peer policy gateways, and no policy gateway is a member of multiple virtual gateways. Then, within AD X, the detected partition will result in the following VGP message exchanges:

- *P* policy gateways each receive at most *P*-1 PG CONNECT messages.

Each policy gateway detecting the adjacent domain partition generates a PG CONNECT message and distributes it to each reachable peer in the virtual gateway.

- P * (V-1) policy gateways each receive at most one VG CONNECT message. The lowest-numbered operational policy gateway in the virtual gateway detecting the partition of the adjacent domain generates a VG CONNECT message and distributes it to a VG representative in all other virtual gateways connected to the domain. In turn, each VG representative distributes the VG CONNECT message to each reachable peer within its virtual gateway.
- P * (V-1) policy gateways each receive at most P-1 PG POLICY messages, and only if the domain has more than a single uniform transit policy. Each policy gateway in each virtual gateway generates a PG POLICY message and distributes it to all reachable peers not currently reachable with respect to the given transit policy.
- P * V policy gateways each receive at most V-1 VG POLICY messages, only if the domain has more than a single uniform transit policy. The lowest-numbered operational policy gateway in each virtual gateway generates a VG POLICY message and distributes it to a VG representative in all other virtual gateways containing at least one reachable member not currently reachable with respect to the given transit policy. In turn, each VG representative distributes a VG POLICY message to each peer within its virtual gateway.

3.5. VGP Message Formats

The virtual gateway protocol number is equal to 0. We describe the contents of each type of VGP message below.

3.5.1. UP/DOWN

The UP/DOWN message type is equal to 0.

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+								
SRC										CMP										DST										AD									
DST										PG										PERIOD										STATE									

SRC CMP

(16 bits) Numeric identifier for the domain component containing the issuing policy gateway.

DST AD (16 bits) Numeric identifier for the destination domain.

DST PG (16 bits) Numeric identifier for the destination policy gateway.

PERIOD (8 bits) Length of the UP/DOWN message generation period, in seconds.

STATE (8 bits) Perceived state (1 up, 0 down) of the direct connection from the perspective of the issuing policy gateway, contained in the right-most bit.

3.5.2. PG CONNECT

The PG CONNECT message type is equal to 1. PG CONNECT messages are not required for any virtual gateway containing exactly two policy gateways.

0										1										2										3																																							
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9																														
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										ADJ AD																				VG																				RQST																			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										NUM RCH																				NUM UNRCH																																							
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
For each reachable adjacent policy gateway:																																																																					
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										ADJ PG																				ADJ CMP																																							
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
For each unreachable adjacent policy gateway:																																																																					
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										ADJ PG																																																											
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							

ADJ AD
(16 bits) Numeric identifier for the adjacent domain.

VG (8 bits) Numeric identifier for the virtual gateway.

RQST (8 bits) Request for a PG CONNECT message (1 request, 0 no request) from each recipient peer, contained in the right-most bit.

NUM RCH (16 bits) Number of adjacent policy gateways within the virtual gateway, which are directly-connected to and currently reachable from the issuing policy gateway.

NUM UNRCH (16 bits) Number of adjacent policy gateways within the


virtual gateway, which are directly-connected to but not currently reachable from the issuing policy gateway.

ADJ PG (16 bits) Numeric identifier for a directly-connected adjacent policy gateway.

ADJ CMP (16 bits) Numeric identifier for the domain component containing the reachable, directly-connected adjacent policy gateway.

3.5.3. PG POLICY

The PG POLICY message type is equal to 2. PG POLICY messages are not required for any virtual gateway containing exactly two policy gateways or for any domain with a single uniform transit policy.



ADJ AD
(16 bits) Numeric identifier for the adjacent domain.

VG (8 bits) Numeric identifier for the virtual gateway.

RQST (8 bits) Request for a PG POLICY message (1 request, 0 no request) from each recipient peer, contained in the right-most bit.

NUM TP (8 bits) Number of transit policies configured to include the virtual gateway.

TP (16 bits) Numeric identifier for a transit policy associated with the virtual gateway.

NUM VG (16 bits) Number of virtual gateways reachable from the issuing policy gateway, via intra-domain routes supporting the transit policy.


UNUSED (8 bits) Not currently used; must be set equal to 0.

NUM CMP (16 bits) Number of adjacent domain components reachable via direct connections through the virtual gateway.

ADJ CMP (16 bits) Numeric identifier for a reachable adjacent domain component.

3.5.4. VG CONNECT

The VG CONNECT message type is equal to 3.



ADJ AD
(16 bits) Numeric identifier for the adjacent domain.

VG (8 bits) Numeric identifier for the virtual gateway.

RQST (8 bits) Request for a VG CONNECT message (1 request, 0 no request) from a recipient in each virtual gateway, contained in the right-most bit.

NUM PG (16 bits) Number of mutually-reachable peer policy gateways in the virtual gateway.


PG (16 bits) Numeric identifier for a peer policy gateway.

NUM CMP (16 bits) Number of components of the adjacent domain reachable via direct connections from the policy gateway.

ADJ CMP (16 bits) Numeric identifier for a reachable adjacent domain component.

3.5.5. VG POLICY

The VG POLICY message type is equal to 4. VG POLICY messages are not required for any domain with a single uniform transit policy.



ADJ AD

(16 bits) Numeric identifier for the adjacent domain.

VG (8 bits) Numeric identifier for the virtual gateway.

RQST (8 bits) Request for a VG POLICY message (1 request, 0 no request) from a recipient in each virtual gateway, contained in the right-most bit.

NUM TP (16 bits) Number of transit policies configured to include the virtual gateway.

TP (16 bits) Numeric identifier for a transit policy associated with the virtual gateway.

NUM GRP (16 bits) Number of groups of virtual gateways, such that all members in a group are reachable from the issuing virtual gateway via intra-domain routes supporting the given transit policy.

NUM VG (16 bits) Number of virtual gateways in a virtual gateway group.

UNUSED (8 bits) Not currently used; must be set equal to 0.

NUM CMP (16 bits) Number of adjacent domain components reachable via direct connections through the virtual gateway.

ADJ CMP (16 bits) Numeric identifier for a reachable adjacent domain component.

Normally, each VG POLICY message will contain a single virtual gateway group. However, if the issuing virtual gateway becomes partitioned such that peers are mutually reachable with respect to some transit policies but not others, virtual gateway groups may be necessary. For example, let PG X and PG Y be two peers in VG A which configured to support transit policies 1 and 2. Suppose that PG X and PG Y are reachable with respect to transit policy 1 but not with respect to transit policy 2. Furthermore, suppose that PG X can reach members of VG B via intra-domain routes of transit policy 2 and that PG Y can reach members of VG C via intra-domain routes of transit policy 2. Then the entry in the VG POLICY message issued by VG A will include, for transit policy 2, two groups of virtual gateways, one containing VG B and one containing VG C.

3.5.6. Negative Acknowledgements

When a policy gateway receives an unacceptable VGP message that passes the CMTTP validation checks, it includes, in its CMTTP ACK, an appropriate VGP negative acknowledgement. This information is placed in the INFORM field of the CMTTP ACK (described previously in section 2.4); the numeric identifier for each type of VGP negative acknowledgement is contained in the left-most 8 bits of the INFORM field. Negative acknowledgements associated with VGP include the following types:

1. Unrecognized VGP message type. Numeric identifier for the unrecognized message type (8 bits).
2. Out-of-date VGP message.
3. Unrecognized virtual gateway source. Numeric identifier for the unrecognized virtual gateway including the adjacent domain identifier (16 bits) and the local identifier (8 bits).

4. Routing Information Distribution

Each domain participating in IDPR generates and distributes its routing information messages to route servers throughout an internetwork. IDPR routing information messages contain information about the transit policies in effect across the given domain and the virtual gateway connectivity to adjacent domains. Route servers in turn use IDPR routing information to generate policy routes between source and destination domains.

There are three different procedures for distributing IDPR routing information:

- The flooding protocol. In this case, a representative policy gateway in each domain floods its routing information messages to route servers in all other domains.
- Remote route server communication. In this case, a route server distributes its domain's routing information messages to route servers in specific destination domains, by encapsulating these messages within IDPR data messages. Thus, a domain administrator may control the distribution of the domain's routing information by restricting routing information exchange with remote route servers. Currently, routing information distribution restrictions are not included in IDPR configuration information.
- The route server query protocol. In this case, a policy gateway or route server requests routing information from another route server, which in turn responds with routing information from its database. The route server query protocol may be used for quickly updating the routing information maintained by a policy gateway or route server that has just been connected or reconnected to an internetwork. It may also be used to retrieve routing information from domains that restrict distribution of their routing information.

In this section, we describe the flooding protocol only. In section 5, we describe the route server query protocol, and in section 5.2, we describe communication between route servers in separate domains.

Policy gateways and route servers use CMTTP for reliable transport of IDPR routing information messages flooded between peer, neighbor, and adjacent policy gateways and between those policy gateways and route servers. The issuing policy gateway must communicate to CMTTP the maximum number of transmissions per routing information message, `flood_ret`, and the interval between routing information message retransmissions, `flood_int` microseconds. The recipient policy gateway or route server must determine routing information message

acceptability, as we describe in section 4.2.3 below.

4.1. AD Representatives

We designate a single policy gateway, the "AD representative", for generating and distributing IDPR routing information about its domain, to ensure that the routing information distributed is consistent and unambiguous and to minimize the communication required for routing information distribution. There is usually only a single AD representative per domain, namely the lowest-numbered operational policy gateway in the domain. Within a domain, policy gateways need no explicit election procedure to determine the AD representative. Instead, all members of a set of policy gateways mutually reachable via intra-domain routes can agree on set membership and therefore on which member has the lowest number.

A partitioned domain has as many AD representatives as it does domain components. In fact, the numeric identifier for an AD representative is identical to the numeric identifier for a domain component. One cannot normally predict when and where a domain partition will occur, and thus any policy gateway within a domain may become an AD representative at any time. To prepare for the role of AD representative in the event of a domain partition, every policy gateway must continually monitor its domain's IDPR routing information, through VGP and through the intra-domain routing procedure.

4.2. Flooding Protocol

An AD representative policy gateway uses unrestricted flooding among all domains to distribute its domain's IDPR routing information messages to route servers in an internetwork. There are two kinds of IDPR routing information messages issued by each AD representative: CONFIGURATION and DYNAMIC messages. Each CONFIGURATION message contains the transit policy information configured by the domain administrator, including for each transit policy, its identifier, its specification, and the sets of virtual gateways configured as mutually reachable via intra-domain routes supporting the given transit policy. Each DYNAMIC message contains information about current virtual gateway connectivity to adjacent domains and about the sets of virtual gateways currently mutually reachable via intra-domain routes supporting the configured transit policies.

The IDPR Flooding Protocol is similar to the flooding procedures described in [9]-[11]. Through flooding, the AD representative distributes its routing information messages to route servers in its own domain and in adjacent domains. After generating a routing information message, the AD representative distributes a copy to each

of its peers and to a selected VG representative (see section 3.1.4) in all other virtual gateways connected to the domain. Each VG representative in turn distributes a copy of the routing information message to each of its peers. We note that distribution of routing information messages among virtual gateways and among peers within a virtual gateway is identical to distribution of inter-VG messages in VGP, as described in section 3.1.3.

Within a virtual gateway, each policy gateway distributes a copy of the routing information message:

- To each route server in its configured set of route servers. A domain administrator should ensure that each route server not contained within a policy gateway appears in the set of configured route servers for at least two distinct policy gateways. Hence, such a route server will continue to receive routing information messages, even when one of the policy gateways becomes unreachable. However, the route server will normally receive duplicate copies of a routing information message.
- To certain directly-connected adjacent policy gateways. A policy gateway distributes a routing information message to a directly-connected adjacent policy gateway in an adjacent domain component, only when it is the lowest-numbered operational peer with a direct connection to the given domain component. We note that each policy gateway knows, through information provided by VGP, which peers have direct connections to which components of the adjacent domain. If the policy gateway has direct connections to more than one adjacent policy gateway in that domain component, it selects the routing information message recipient according to the order in which the adjacent policy gateways appear in its database, choosing the first one encountered. This selection procedure ensures that a copy of the routing information message reaches each component of the adjacent domain, while limiting the number of copies distributed.

Once a routing information message reaches an adjacent policy gateway, that policy gateway distributes copies of the message throughout its domain. The adjacent policy gateway, acting as the first recipient of the routing information message in its domain, follows the same message distribution procedure as the AD representative in the source domain, as described above. The flooding procedure terminates when all reachable route servers in an internetwork receive a copy of the routing information message.

Neighbor policy gateways may receive copies of the same routing information message from different adjoining domains. If two neighbor policy gateways receive the message copies simultaneously,

they will distribute them to VG representatives in other virtual gateways within their domain, resulting in duplicate message distribution. However, each policy gateway stops the spread of duplicate routing information messages as soon as it detects them, as described in section 4.2.3 below. In the Internet, we expect simultaneous message receptions to be the exception rather than the rule, given the hierarchical structure of the current topology.

4.2.1. Message Generation

An AD representative generates and distributes a CONFIGURATION message whenever there is a configuration change in a transit policy or virtual gateway associated with its domain. This ensures that route servers maintain an up-to-date view of a domain's configured transit policies and adjacencies. An AD representative may also distribute a CONFIGURATION message at a configurable period of `conf_per` (500) hours. A CONFIGURATION message contains, for each configured transit policy, the identifier assigned by the domain administrator, the specification, and the set of associated "virtual gateway groups". Each virtual gateway group comprises virtual gateways configured to be mutually reachable via intra-domain routes of the given transit policy. Accompanying each virtual gateway listed is an indication of whether the virtual gateway is configured to be a domain entry point, a domain exit point, or both according to the given transit policy. The CONFIGURATION message also contains the set of local route servers that the domain administrator has configured to be available to IDPR clients in other domains.

An AD representative generates and distributes a DYNAMIC message whenever there is a change in currently supported transit policies or in current virtual gateway connectivity associated with its domain. This ensures that route servers maintain an up-to-date view of a domain's supported transit policies and existing adjacencies and how they differ from those configured for the domain. Specifically, an AD representative generates a DYNAMIC message whenever there is a change in the connectivity, through the given domain and with respect to a configured transit policy, between two components of adjoining domains. An AD representative may also distribute a DYNAMIC message at a configurable period of `dyn_per` (24) hours. A DYNAMIC message contains, for each configured transit policy, its identifier, associated virtual gateway groups, and domain components reachable through virtual gateways in each group. Each DYNAMIC message also contains the set of currently "unavailable", either down or unreachable, virtual gateways in the domain.

We note that each virtual gateway group expressed in a DYNAMIC message may be a proper subset of one of the corresponding virtual gateway groups expressed in a CONFIGURATION message. For example,

suppose that, for a given domain, the virtual gateway group (VG A,...,VG E) were configured for a transit policy such that each virtual gateway was both a domain entry and exit point. Thus, all virtual gateways in this group are configured to be mutually reachable via intra-domain routes of the given transit policy. Now suppose that VG E becomes unreachable because of a power failure and furthermore that the remaining virtual gateways form two distinct groups, (VG A,VG B) and (VG C,VG D), such that although virtual gateways in both groups are still mutually reachable via some intra-domain routes they are no longer mutually reachable via any intra-domain routes of the given transit policy. In this case, the virtual gateway groups for the given transit policy now become (VG A,VG B) and (VG C,VG D); VG E is listed as an unavailable virtual gateway.

A route server uses information about the set of unavailable virtual gateways to determine which of its routes are no longer viable, and it subsequently removes such routes from its route database. Although route servers could determine the set of unavailable virtual gateways using information about configured virtual gateways and currently reachable virtual gateways, the associated processing cost is high. In particular, a route server would have to examine all virtual gateway groups listed in a DYNAMIC message to determine whether there are any unavailable virtual gateways in the given domain. To reduce the message processing at each route server, we have chosen to include the set of unavailable virtual gateways in each DYNAMIC message.

In order to construct a DYNAMIC message, an AD representative assembles information gathered from intra-domain routing and from VGP. Specifically, the AD representative uses the following information:

- VG CONNECT and UP/DOWN messages to determine the state, up or down, of each of its domain's virtual gateways and the adjacent domain components reachable through a given virtual gateway.
- Intra-domain routing information to determine, for each of its domain's transit policies, whether a given virtual gateway in the domain is reachable with respect to that transit policy.
- VG POLICY messages to determine the connectivity of adjoining domain components, across the given domain and with respect to a configured transit policy, such that these components are adjacent to virtual gateways not currently reachable from the AD representative's virtual gateway according to the given transit policy.

4.2.2. Sequence Numbers

Each IDPR routing information message carries a sequence number which, when used in conjunction with the timestamp carried in the CMTTP message header, determines the recency of the message. An AD representative assigns a sequence number to each routing information message it generates, depending upon its internal clock time:

- The AD representative sets the sequence number to 0, if its internal clock time is greater than the timestamp in its previously generated routing information message.
- The AD representative sets the sequence number to 1 greater than the sequence number in its previously generated routing information message, if its internal clock time equals the timestamp for its previously generated routing information message and if the previous sequence number is less than the maximum value (currently $2^{16} - 1$). If the previous sequence number equals the maximum value, the AD representative waits until its internal clock time exceeds the timestamp in its previously generated routing information message and then sets the sequence number to 0.

In general, we do not expect generation of multiple distinct IDPR routing information messages carrying identical timestamps, and so the sequence number may seem superfluous. However, the sequence number may become necessary during synchronization of an AD representative's internal clock. In particular, the AD representative may need to freeze the clock value during synchronization, and thus distinct sequence numbers serve to distinguish routing information messages generated during the clock synchronization interval.

4.2.3. Message Acceptance

Prior to a policy gateway forwarding a routing information message or a route server incorporating routing information into its routing information database, the policy gateway or route server assesses routing information message acceptability. An IDPR routing information message is "acceptable" if:

- It passes the CMTTP validation checks.
- Its timestamp is less than `conf_old` (530) hours behind the recipient's internal clock time for CONFIGURATION messages and less than `dyn_old` (25) hours behind the recipient's internal clock time for DYNAMIC messages.
- Its timestamp and sequence number indicate that it is more recent

than the currently-stored routing information from the given domain. If there is no routing information currently stored from the given domain, then the routing information message contains, by default, the more recent information.

Policy gateways acknowledge and forward acceptable IDPR routing information messages, according to the flooding protocol described in section 4.2 above. Moreover, each policy gateway retains the timestamp and sequence number for the most recently accepted routing information message from each domain and uses these values to determine acceptability of routing information messages received in the future. Route servers acknowledge the receipt of acceptable routing information messages and incorporate the contents of these messages into their routing information databases, contingent upon criteria discussed in section 4.2.4 below; however, they do not participate in the flooding protocol. We note that when a policy gateway or route server first returns to service, it immediately updates its routing information database with routing information obtained from another route server, using the route server query protocol described in section 5.

An AD representative takes special action upon receiving an acceptable routing information message, supposedly generated by itself but originally obtained from a policy gateway or route server other than itself. There are at least three possible reasons for the occurrence of this event:

- The routing information message has been corrupted in a way that is not detectable by the integrity/authentication value.
- The AD representative has experienced a memory error.
- Some other entity is generating routing information messages on behalf of the AD representative.

In any case, the AD representative logs the event for network management. Moreover, the AD representative must reestablish its own routing information messages as the most recent for its domain. To do so, the AD representative waits until its internal clock time exceeds the value of the timestamp in the received routing information message and then generates a new routing information message using the currently-stored domain routing information supplied by VGP and by the intra-domain routing procedure. Note that the length of time the AD representative must wait to generate the new message is at most `cmtpl_new` (300) seconds, the maximum CMTP-tolerated difference between the received message's timestamp and the AD representative's internal clock time.

IDPR routing information messages that pass the CMTTP validity checks but appear less recent than stored routing information are unacceptable. Policy gateways do not forward unacceptable routing information messages, and route servers do not incorporate the contents of unacceptable routing information messages into their routing information databases. Instead, the recipient of an unacceptable routing information message acknowledges the message in one of two ways:

- If the routing information message timestamp and sequence number equal to the timestamp and sequence number associated with the stored routing information for the given domain, the recipient assumes that the routing information message is a duplicate and acknowledges the message.
- If the routing information message timestamp and sequence number indicate that the message is less recent than the stored routing information for the domain, the recipient acknowledges the message with an indication that the routing information it contains is out-of-date. Such a negative acknowledgement is a signal to the sender of the routing information message to request more up-to-date routing information from a route server, using the route server query protocol. Furthermore, if the recipient of the out-of-date routing information message is a route server, it regenerates a routing information message from its own routing information database and forwards the message to the sender. The sender may in turn propagate this more recent routing information message to other policy gateways and route servers.

4.2.4. Message Incorporation

A route server usually stores the entire contents of an acceptable IDPR routing information message in its routing information database, so that it has access to all advertised transit policies when generating a route and so that it can regenerate routing information messages at a later point in time if requested to do so by another route server or policy gateway. However, a route server may elect not to store all routing information message contents. In particular, the route server need not store any transit policy that excludes the route server's domain as a source or any routing information from a domain that the route server's domain's source policies exclude for transit. Selective storing of routing information message contents simplifies the route generation procedure since it reduces the search space of possible routes, and it limits the amount of route server memory devoted to routing information. However, selective storing of routing information also means that the route server cannot always regenerate the original routing information message, if requested to do so by another route

server or policy gateway.

An acceptable IDPR routing information message may contain transit policy information that is not well-defined according to the route server's perception. A CONFIGURATION message may contain an unrecognized domain, virtual gateway, or transit policy attribute, such as user class access restrictions or offered service. In this case, "unrecognized" means that the value in the routing information message is not listed in the route server's configuration database, as described previously in section 1.8.2. A DYNAMIC message may contain an unrecognized transit policy or virtual gateway. In this case, "unrecognized" means that the transit policy or virtual gateway was not listed in the most recent CONFIGURATION message for the given domain.

Each route server can always parse an acceptable routing information message, even if some of the information is not well-defined, and thus can always use the information that it does recognize. Therefore, a route server can store the contents of acceptable routing information messages from domains in which it is interested, regardless of whether all contents appear to be well-defined at present. If a routing message contains unrecognized information, the route server may attempt to obtain the additional information it needs to decipher the unrecognized information. For a CONFIGURATION message, the route server logs the event for network management; for a DYNAMIC message, the route server requests, from another route server, the most recent CONFIGURATION message for the domain in question.

When a domain is partitioned, each domain component has its own AD representative, which generates routing information messages on behalf of that component. Discovery of a domain partition prompts the AD representative for each domain component to generate and distribute a DYNAMIC message. In this case, a route server receives and stores more than one routing information message at a time for the given domain, namely one for each domain component.

When the partition heals, the AD representative for the entire domain generates and distributes a DYNAMIC message. In each route server's routing information database, the new DYNAMIC message does not automatically replace all of the currently-stored DYNAMIC messages for the given domain. Instead, the new message only replaces that message whose AD representative matches the AD representative for the new message. The other DYNAMIC messages, generated during the period over which the partition occurred, remain in the routing information database until they attain their maximum lifetime, as described in section 4.2.5 below. Such stale information may lead to the generation of routes that result in path setup failures and hence the

selection of alternative routes. To reduce the chances of path setup failures, we will investigate, for a future version of IDPR, mechanisms for removing partition-related DYNAMIC messages immediately after a partition disappears.

4.2.5. Routing Information Database

We expect that most of the IDPR routing information stored in a routing information database will remain viable for long periods of time, perhaps until a domain reconfiguration occurs. By "viable", we mean that the information reflects the current state of the domain and hence may be used successfully for generating policy routes. To reduce the probability of retaining stale routing information, a route server imposes a maximum lifetime on each database entry, initialized when it incorporates an accepted entry into its routing information database. The maximum lifetime should be longer than the corresponding message generation period, so that the database entry is likely to be refreshed before it attains its maximum lifetime.

Each CONFIGURATION message stored in the routing information database has a maximum lifetime of `conf_old` (530) hours; each DYNAMIC message stored in the routing information database has a maximum lifetime of `dyn_old` (25) hours. Periodic generation of routing information messages makes it unlikely that any routing information message will remain in a routing information database for its full lifetime. However, a routing information message may attain its maximum lifetime in a route server that is separated from a internetwork for a long period of time.

When an IDPR routing information message attains its maximum lifetime in a routing information database, the route server removes the message contents from its database, so that it will not generate new routes with the outdated routing information nor distribute old routing information in response to requests from other route servers or policy gateways. Nevertheless, the route server continues to dispense routes previously generated with the old routing information, as long as path setup (see section 7) for these routes succeeds.

The route server treats routing information message lifetime expiration differently, depending on the type of routing information message. When a CONFIGURATION message expires, the route server requests, from another route server, the most recent CONFIGURATION message issued for the given domain. When a DYNAMIC message expires, the route server does not initially attempt to obtain more recent routing information. Instead, if route generation is necessary, the route server uses the routing information contained in the corresponding CONFIGURATION message for the given domain. Only if

there is a path setup failure (see section 7.4) involving the given domain does the route server request, from another route server, the most recent DYNAMIC message issued for the given domain.

4.3. Routing Information Message Formats

The flooding protocol number is equal to 1. We describe the contents of each type of routing information message below.

4.3.1. CONFIGURATION

The CONFIGURATION message type is equal to 0.

0										1										2										3											
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1
AD CMP										SEQ																															
NUM TP										NUM RS																															
RS																																									
For each transit policy configured for the domain:																																									
TP										NUM ATR																															
For each attribute of the transit policy:																																									
ATR TYP										ATR LEN																															
For the source/destination access restrictions attribute:																																									
NUM AD GRP																																									
For each domain group in the source/destination access restrictions:																																									
NUM AD										AD																															
AD FLGS					NUM HST					HST SET																															
For the temporal access restrictions attribute:																																									
NUM TIM																																									

For each set of times in the temporal access restrictions:

+-----+	
TIM FLGS	DURATION
+-----+	
START	
+-----+	
PERIOD	ACTIVE
+-----+	

For the user class access restrictions attribute:

+-----+	
NUM UCI	
+-----+	

For each UCI in the user class access restrictions:

+-----+	
UCI	
+-----+	

For each offered service attribute:

+-----+	
OFR SRV	
+-----+	

For the virtual gateway access restrictions attribute:

+-----+	
NUM VG GRP	
+-----+	

For each virtual gateway group in the virtual gateway access restrictions:

+-----+			
NUM VG	ADJ AD		
+-----+			
VG	VG FLGS		
+-----+			

AD CMP

(16 bits) Numeric identifier for the domain component containing the AD representative policy gateway.

SEQ (16 bits) Routing information message sequence number.

NUM TP (16 bits) Number of transit policy specifications contained in the routing information message.

NUM RS (16 bits) Number of route servers advertised to serve clients outside of the domain.

RS (16 bits) Numeric identifier for a route server.

TP (16 bits) Numeric identifier for a transit policy specification.

NUM ATR (16 bits) Number of attributes associated with the transit policy.

ATR TYP (16 bits) Numeric identifier for a type of attribute. Valid attributes include the following types:

- Set of virtual gateway access restrictions (see section 1.4.2) associated with the transit policy (variable). This attribute must be included.
- Set of source/destination access restrictions (see section 1.4.2) associated with the transit policy (variable). This attribute may be omitted. Absence of this attribute implies that traffic from any source to any destination is acceptable.
- Set of temporal access restrictions (see section 1.4.2) associated with the transit policy (variable). This attribute may be omitted. Absence of this attribute implies that the transit policy applies at all times.
- Set of user class access restrictions (see section 1.4.2) associated with the transit policy (variable). This attribute may be omitted. Absence of this attribute implies that traffic from any user class is acceptable.
- Average delay in milliseconds (16 bits). This attribute may be omitted.
- Delay variation in milliseconds (16 bits). This attribute may be omitted.
- Average available bandwidth in bits per second (48 bits). This attribute may be omitted.
- Available bandwidth variation in bits per second (48 bits). This attribute may be omitted.
- MTU in bytes (16 bits). This attribute may be omitted.
- Charge per byte in thousandths of a cent (16 bits). This attribute may be omitted.
- Charge per message in thousandths of a cent (16 bits). This attribute may be omitted.
- Charge for session time in thousandths of a cent per second (16 bits). This attribute may be omitted. Absence of any charge attribute implies that the domain provides free transit service.

ATR LEN (16 bits) Length of an attribute in bytes, beginning with the subsequent field.

NUM AD GRP (16 bits) Number of source/destination domain groups (see section 1.4.2) associated with the source/destination access restrictions.

NUM AD (16 bits) Number of domains or sets of domains in a domain group.

AD (16 bits) Numeric identifier for a domain or domain set.

AD FLGS (8 bits) Set of five flags indicating how to interpret the AD field, contained in the right-most bits. Proceeding left to right, the first flag indicates whether the transit policy applies to all domains or to specific domains (1 all, 0 specific), and when set to 1, causes the second and third flags to be ignored. The second flag indicates whether the domain identifier signifies a single domain or a domain set (1 single, 0 set). The third flag indicates whether the transit policy applies to the given domain or domain set (1 applies, 0 does not apply) and is used for representing complements of sets of domains. The fourth flag indicates whether the domain is a source (1 source, 0 not source). The fifth flag indicates whether the domain is a destination (1 destination, 0 not destination). At least one of the fourth and fifth flags must be set to 1.

NUM HST (8 bits) Number of "host sets" (see section 1.4.2) associated with a particular domain or domain set. The value 0 indicates that all hosts in the given domain or domain set are acceptable sources or destinations, as specified by the fourth and fifth AD flags.

HST SET (16 bits) Numeric identifier for a host set.

NUM TIM (16 bits) Number of time specifications associated with the temporal access restrictions. Each time specification is split into a set of contiguous identical periods, as we describe below.

TIM FLGS (8 bits) Set of two flags indicating how to combine the time specifications, contained in the right-most bits. Proceeding left to right, the first flag indicates whether the transit policy applies during the periods specified in the time specification (1 applies, 0 does not apply) and is used for representing complements of policy applicability intervals. The second flag indicates whether the time specification takes precedence over the previous time specifications listed (1 precedence, 0 no precedence). Precedence is equivalent to the boolean OR and AND operators, in the following sense. At any given instant, a transit policy either applies or does not apply, according to a given time specification, and we can assign a boolean

value to the state of policy applicability according to a given time specification. If the second flag assumes the value 1 for a given time specification, that indicates the boolean operator OR should be applied to the values of policy applicability, according to the given time specification and to all previously listed time specifications. If the second flag assumes the value 0 for a given time specification, that indicates the boolean operator AND should be applied to the values of policy applicability, according to the given time specification and to all previously listed time specifications.

DURATION (24 bits) Length of the time specification duration, in minutes. A value of 0 indicates an infinite duration.

START (32 bits) Time at which the time specification first takes effect, in seconds elapsed since 1 January 1970 0:00 GMT.

PERIOD (16 bits) Length of each time period within the time specification, in minutes.

ACTIVE (16 bits) Length of the policy applicable interval during each time period, in minutes from the beginning of the time period.

NUM UCI (16 bits) Number of user classes associated with the user class access restrictions.

UCI (8 bits) Numeric identifier for a user class.

NUM VG GRP (16 bits) Number of virtual gateway groups (see section 1.4.2) associated with the virtual gateway access restrictions.

NUM VG (16 bits) Number of virtual gateways in a virtual gateway group.


ADJ AD (16 bits) Numeric identifier for the adjacent domain to which a virtual gateway connects.

VG (8 bits) Numeric identifier for a virtual gateway.

VG FLGS (8 bits) Set of two flags indicating how to interpret the VG field, contained in the right-most bits. Proceeding left to right, the first flag indicates whether the virtual gateway is a domain entry point (1 entry, 0 not entry). The second flag indicates whether the virtual gateway is a domain exit point (1 exit, 0 not exit). At least one of the first and second flags must be set to 1.

4.3.2. DYNAMIC

The DYNAMIC message type is equal to 1.



AD CMP

(16 bits) Numeric identifier for the domain component containing the AD representative policy gateway.

SEQ (16 bits) Routing information message sequence number.

UNAVL VG (16 bits) Number of virtual gateways in the domain component that are currently unavailable via any intra-domain routes.

NUM PS (16 bits) Number of sets of transit policies listed. Transit policy sets provide a mechanism for reducing the size of DYNAMIC messages. A single set of virtual gateway groups applies to all transit policies in a given set.

ADJ AD (16 bits) Numeric identifier for the adjacent domain to which a virtual gateway connects.

VG (8 bits) Numeric identifier for a virtual gateway.

UNUSED (8 bits) Not currently used; must be set equal to 0.

NUM TP (16 bits) Number of transit policies in a set.

NUM VGGRP (16 bits) Number of virtual gateway groups currently associated with the transit policy set.

TP (16 bits) Numeric identifier for a transit policy.

NUM VG (16 bits) Number of virtual gateways in a virtual gateway group.

VG FLGS (8 bits) Set of two flags indicating how to interpret the VG field, contained in the right-most bits. Proceeding left to right, the first flag indicates whether the virtual gateway is a domain entry point (1 entry, 0 not entry). The second flag indicates whether the virtual gateway is a domain exit point (1 exit, 0 not exit). At least one of the first and second flags must be set to 1.

NUM CMP (16 bits) Number of adjacent domain components reachable via direct connections through the virtual gateway.

ADJ CMP (16 bits) Numeric identifier for a reachable adjacent domain component.

4.3.3. Negative Acknowledgements

When a policy gateway or route server receives an unacceptable IDPR routing information message that passes the CMTTP validation checks, it includes, in its CMTTP ACK, an appropriate negative acknowledgement. This information is placed in the INFORM field of the CMTTP ACK (described previously in section 2.4); the numeric identifier for each type of routing information message negative acknowledgement is contained in the left-most 8 bits of the INFORM field. Negative acknowledgements associated with routing information messages include the following types:

1. Unrecognized IDPR routing information message type. Numeric identifier for the unrecognized message type (8 bits).
2. Out-of-date IDPR routing information message. This is a signal to the sender that it may not have the most recent routing information for the given domain.

5. Route Server Query Protocol

Each route server is responsible for maintaining both the routing information database and the route database and for responding to database information requests from policy gateways and other route servers. These requests and their responses are the messages exchanged via the Route Server Query Protocol (RSQP).

Policy gateways and route servers normally invoke RSQP to replace absent, outdated, or corrupted information in their own routing information or route databases. In section 4, we discussed some of the situations in which RSQP may be invoked; in this section and in section 7, we discuss other such situations.

5.1. Message Exchange

Policy gateways and route servers use CMTP for reliable transport of route server requests and responses. RSQP must communicate to CMTP the maximum number of transmissions per request/response message, `rsqp_ret`, and the interval between request/response message retransmissions, `rsqp_int` microseconds. A route server request/response message is "acceptable" if:

- It passes the CMTP validation checks.
- Its timestamp is less than `rsqp_old` (300) seconds behind the recipient's internal clock time.

With RSQP, a requesting entity expects to receive an acknowledgement from the queried route server indicating whether the route server can accommodate the request. The route server may fail to fill a given request for either of the following reasons:

- Its corresponding database contains no entry or only a partial entry for the requested information.
- It is governed by special message distribution rules, imposed by the domain administrator, that preclude it from releasing the requested information. Currently, such distribution rules are not included in IDPR configuration information.

For all requests that it cannot fill, the route server responds with a negative acknowledgement message carried in a CMTP acknowledgement, indicating the set of unfulfilled requests (see section 5.5.4).

If the requesting entity either receives a negative acknowledgement or does not receive any acknowledgement after `rsqp_ret` attempts directed at the same route server, it queries a different route

server, as long as the number of attempted requests to different route servers does not exceed `rsqp_try` (3). Specifically, the requesting entity proceeds in round-robin order through its list of addressable route servers. However, if the requesting entity is unsuccessful after `rsqp_try` attempts, it abandons the request altogether and logs the event for network management.

A policy gateway or a route server can request information from any route server that it can address. Addresses for local route servers within a domain are part of the configuration for each IDPR entity within a domain; addresses for remote route servers in other domains are obtained through flooded CONFIGURATION messages, as described previously in section 4.2.1. However, requesting entities always query local route servers before remote route servers, in order to contain the costs associated with the query and response. If the requesting entity and the queried route server are in the same domain, they can communicate over intra-domain routes, whereas if the requesting entity and the queried route server are in different domains, they must obtain a policy route and establish a path before they can communicate, as we describe below.

5.2. Remote Route Server Communication

RSQP communication involving a remote route server requires a policy route and accompanying path setup (see section 7) between the requesting and queried entities, as these entities reside in different domains. After generating a request message, the requesting entity hands to CMTP its request message along with the remote route server's entity and domain identifiers. CMTP encloses the request in a DATAGRAM and hands the DATAGRAM and remote route server information to the path agent. Using the remote route server information, the path agent obtains, and if necessary sets up, a path to the remote route server. Once the path to the remote route server has been successfully established, the path agent encapsulates the DATAGRAM within an IDPR data message and forwards the data message along the designated path.

When the path agent in the remote route server receives the IDPR data message, it extracts the DATAGRAM and hands it to CMTP. In addition, the path agent, using the requesting entity and domain identifiers contained in the path identifier, obtains, and if necessary sets up, a path back to the requesting entity.

If the DATAGRAM fails any of the CMTP validation checks, CMTP returns a NAK to the requesting entity. If the DATAGRAM passes all of the CMTP validation checks, the remote route server assesses the acceptability of the request message. Provided the request message is acceptable, the remote route server determines whether it can

fulfill the request and directs CMTTP to return an ACK to the requesting entity. The ACK may contain a negative acknowledgement if the entire request cannot be fulfilled.

The remote route server generates responses for all requests that it can fulfill and returns the responses to the requesting entity. Specifically, the remote route server hands to CMTTP its response and the requesting entity information. CMTTP in turn encloses the response in a DATAGRAM.

When returning an ACK, a NAK, or a response to the requesting entity, the remote route server hands the corresponding CMTTP message and requesting entity information to the path agent. Using the requesting entity information, the path agent retrieves the path to the requesting entity, encapsulates the CMTTP message within an IDPR data message, and forwards the data message along the designated path.

When the path agent in the requesting entity receives the IDPR data message, it extracts the ACK, NAK, or response to its request and performs the CMTTP validation checks for that message. In the case of a response message, the requesting entity also assesses message acceptability before incorporating the contents into the appropriate database.

5.3 Routing Information

Policy gateways and route servers request routing information from route servers, in order to update their routing information databases. To obtain routing information from a route server, the requesting entity issues a ROUTING INFORMATION REQUEST message containing the type of routing information requested - CONFIGURATION messages, DYNAMIC messages, or both - and the set of domains from which the routing information is requested.

Upon receiving a ROUTING INFORMATION REQUEST message, a route server first assesses message acceptability before proceeding to act on the contents. If the ROUTING INFORMATION REQUEST message is deemed acceptable, the route server determines how much of the request it can fulfill and then instructs CMTTP to generate an acknowledgement, indicating its ability to fulfill the request. The route server proceeds to fulfill as much of the request as possible by reconstructing individual routing information messages, one per requested message type and domain, from its routing information database. We note that only a regenerated routing information message whose entire contents match that of the original routing information message may pass the CMTTP integrity/authentication checks.

5.4. Routes

Path agents request routes from route servers when they require policy routes for path setup. To obtain routes from a route server, the requesting path agent issues a ROUTE REQUEST message containing the destination domain and applicable service requirements, the maximum number of routes requested, a directive indicating whether to generate the routes or retrieve them from the route database, and a directive indicating whether to refresh the routing information database with the most recent CONFIGURATION or DYNAMIC message from a given domain, before generating the routes. To refresh its routing information database, a route server must obtain routing information from another route server. The path agent usually issues routing information database refresh directives in response to a failed path setup. We discuss the application of these directives in more detail in section 7.4.


Upon receiving a ROUTE REQUEST message, a route server first assesses message acceptability before proceeding to act on the contents. If the ROUTE REQUEST message is deemed acceptable, the route server determines whether it can fulfill the request and then instructs CMTDP to generate an acknowledgement, indicating its ability to fulfill the request. The route server proceeds to fulfill the request with policy routes, either retrieved from its route database or generated from its routing information database if necessary, and returns these routes in a ROUTE RESPONSE message.

5.5. Route Server Message Formats

The route server query protocol number is equal to 2. We describe the contents of each type of RSQP message below.

5.5.1. ROUTING INFORMATION REQUEST

The ROUTING INFORMATION REQUEST message type is equal to 0.



QRY AD

(16 bits) Numeric identifier for the domain containing the

queried route server.

QRY RS (16 bits) Numeric identifier for the queried route server.

NUM AD (16 bits) Number of domains about which routing information is requested. The value 0 indicates a request for routing information from all domains.

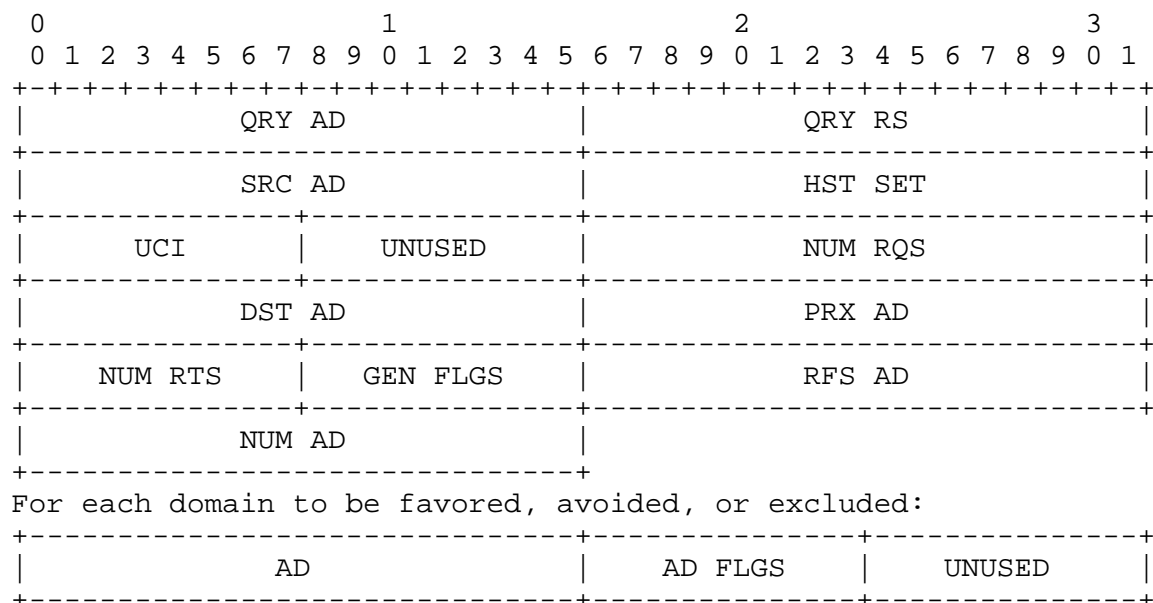
AD (16 bits) Numeric identifier for a domain. This field is absent when NUM AD equals 0.

RIM FLGS (8 bits) Set of two flags indicating the type of routing information messages requested, contained in the right-most bits. Proceeding left to right, the first flag indicates whether the request is for a CONFIGURATION message (1 CONFIGURATION, 0 no CONFIGURATION). The second flag indicates whether the request is for a DYNAMIC message (1 DYNAMIC, 0 no DYNAMIC). At least one of the first and second flags must be set to 1.

UNUSED (8 bits) Not currently used; must be set equal to 0.

5.5.2. ROUTE REQUEST

The ROUTE REQUEST message type is equal to 1.



For each requested service:

+-----+-----+	
RQS TYP	RQS LEN
+-----+-----+	
RQS SRV	
+-----+-----+	

QRY AD

(16 bits) Numeric identifier for the domain containing the queried route server.

QRY RS (16 bits) Numeric identifier for the queried route server.

SRC AD (16 bits) Numeric identifier for the route's source domain.

HST SET (16 bits) Numeric identifier for the source's host set.

UCI (8 bits) Numeric identifier for the source user class. The value 0 indicates that there is no particular source user class.

UNUSED (8 bits) Not currently used; must be set equal to 0.

NUM RQS (16 bits) Number of requested services. The value 0 indicates that the source requests no special services.

DST AD (16 bits) Numeric identifier for the route's destination domain.

PRX AD (16 bits) Numeric identifier for the destination domain's proxy (see section 1.3.1). If the destination domain provides the path agent function for its hosts, then the destination and proxy domains are identical. A route server constructs routes between the source domain's proxy and the destination domain's proxy. We note that the source domain's proxy is identical to the domain issuing the CMTTP message containing the ROUTE REQUEST message, and hence available in the CMTTP header.

NUM RTS (8 bits) Number of policy routes requested.

GEN FLGS (8 bits) Set of three flags indicating how to obtain the requested routes, contained in the right-most bits. Proceeding left to right, the first flag indicates whether the route server should retrieve existing routes from its route database or generate new routes (1 retrieve, 0 generate). The second flag indicates whether the route server should refresh its routing information database before generating the requested routes (1 refresh, 0 no refresh) and when set to 1, causes the third flag and the RFS AD field to become significant. The third flag

indicates whether the routing information database refresh should include CONFIGURATION messages or DYNAMIC messages (1 configuration, 0 dynamic).

RFS AD (16 bits) Numeric identifier for the domain for which routing information should be refreshed. This field is meaningful only if the second flag in the GEN FLGS field is set to 1.

NUM AD (16 bits) Number of transit domains that are to be favored, avoided, or excluded during route selection (see section 1.4.1).

AD (16 bits) Numeric identifier for a transit domain to be favored, avoided, or excluded.

AD FLGS (8 bits) Three flags indicating how to interpret the AD field, contained in the right-most bits. Proceeding left to right, the first flag indicates whether the domain should be favored (1 favored, 0 not favored). The second flag indicates whether the domain should be avoided (1 avoided, 0 not avoided). The third flag indicates whether the domain should be excluded (1 excluded, 0 not excluded). No more than one of the first, second, and third flags must set to 1.

RQS TYP (16 bits) Numeric identifier for a type of requested service. Valid requested services include the following types:

1. Upper bound on delay, in milliseconds (16 bits). This attribute may be omitted.
2. Minimum delay route. This attribute may be omitted.
3. Upper bound on delay variation, in milliseconds (16 bits). This attribute may be omitted.
4. Minimum delay variation route. This attribute may be omitted.
5. Lower bound on bandwidth, in bits per second (48 bits). This attribute may be omitted.
6. Maximum bandwidth route. This attribute may be omitted.
7. Upper bound on monetary cost, in cents (32 bits). This attribute may be omitted.
8. Minimum monetary cost route. This attribute may be omitted.
9. Path lifetime in minutes (16 bits). This attribute may be omitted but must be present if types 7 or 8 are present. Route servers

use path lifetime information together with domain charging method to compute expected session monetary cost over a given domain.

10. Path lifetime in messages (16 bits). This attribute may be omitted but must be present if types 7 or 8 are present.
11. Path lifetime in bytes (48 bits). This attribute may be omitted but must be present if types 7 or 8 are present.

RQS LEN

(16 bits) Length of the requested service, in bytes, beginning with the next field.

RQS SRV

(variable) Description of the requested service.

5.5.3. ROUTE RESPONSE

The ROUTE RESPONSE message type is equal to 2.

0																1																2																3															
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																																
+-----+																+-----+																+-----+																+-----+															
NUM RTS																																																															
+-----+																																																															

For each route provided:

+-----+																+-----+															
NUM AD																RTE FLGS															
+-----+																+-----+															

For each domain in the route:

+-----+																+-----+																+-----+																+-----+															
AD LEN																VG																ADJ AD																															
+-----+																+-----+																+-----+																+-----+															
																ADJ CMP																NUM TP																															
+-----+																+-----+																+-----+																+-----+															
																TP																																															
+-----+																+-----+																+-----+																+-----+															

NUM RTS

(16 bits) Number of policy routes provided.

RTE FLGS (8 bits) Set of two flags indicating the directions in which a route can be used, contained in the right-most bits. Refer to sections 6.2, 7, and 7.2 for detailed discussions of path directionality. Proceeding left to right, the first flag indicates whether the route can be used from source to destination (1 from source, 0 not from source). The second flag

indicates whether the route can be used from destination to source (1 from destination, 0 not from destination). At least one of the first and second flags must be set to 1, if NUM RTS is greater than 0.

NUM AD (8 bits) Number of domains in the policy route, not including the first domain on the route.

AD LEN (8 bits) Length of the information associated with a particular domain, in bytes, beginning with the next field.

VG (8 bits) Numeric identifier for an exit virtual gateway.

ADJ AD (16 bits) Numeric identifier for the adjacent domain connected to the virtual gateway.

ADJ CMP (16 bits) Numeric identifier for the adjacent domain component. Used by policy gateways to select a route across a virtual gateway connecting to a partitioned domain.

NUM TP (16 bits) Number of transit policies that apply to the section of the route traversing the domain component.

TP (16 bits) Numeric identifier for a transit policy.

5.5.4. Negative Acknowledgements

When a policy gateway receives an unacceptable RSQP message that passes the CMTTP validation checks, it includes, in its CMTTP ACK, an appropriate negative acknowledgement. This information is placed in the INFORM field of the CMTTP ACK (described previously in section 2.4); the numeric identifier for each type of RSQP negative acknowledgement is contained in the left-most 8 bits of the INFORM field. Negative acknowledgements associated with RSQP include the following types:

1. Unrecognized RSQP message type. Numeric identifier for the unrecognized message type (8 bits).
2. Out-of-date RSQP message.
3. Unable to fill requests for routing information from the following domains. Number of domains for which requests cannot be filled (16 bits); a value of 0 indicates that the route server cannot fill any of the requests. Numeric identifier for each domain for which a request cannot be filled (16 bits).

4. Unable to fill requests for routes to the following destination domain. Numeric identifier for the destination domain (16 bits).

6. Route Generation

Route generation is the most computationally complex part of IDPR, because of the number of domains and the number and heterogeneity of policies that it must accommodate. Route servers must generate policy routes that satisfy the requested services of the source domains and respect the offered services of the transit domains.

We distinguish requested qualities of service and route generation with respect to them as follows:

- Requested service limits include upper bounds on route delay, route delay variation, and session monetary cost and lower bounds on available route bandwidth. Generating a route that must satisfy more than one quality of service constraint, for example route delay of no more than X seconds and available route bandwidth of no less than Y bits per second, is an NP-complete problem.
- Optimal requested services include minimum route delay, minimum route delay variation, minimum session monetary cost, and maximum available route bandwidth. In the worst case, the computational complexity of generating a route that is optimal with respect to a given requested service is $O((N + L) \log N)$ for Dijkstra's shortest path first (SPF) search and $O(N + (L * L))$ for breadth-first (BF) search, where N is the number of nodes and L is the number of links in the search graph. Multi-criteria optimization, for example finding a route with minimal delay variation and minimal session monetary cost, may be defined in several ways. One approach to multi-criteria optimization is to assign each link a single value equal to a weighted sum of the values of the individual offered qualities of service and generate a route that is optimal with respect to this new criterion. However, selecting the weights that yield the desired route generation behavior is itself an optimization procedure and hence not trivial.

To help contain the combinatorial explosion of processing and memory costs associated with route generation, we supply the following guidelines for generation of suitable policy routes:

- Each route server should only generate policy routes from the perspective of its own domain as source; it need not generate policy routes for arbitrary source/destination domain pairs. Thus, we can distribute the computational burden over all route servers.
- Route servers should precompute routes for which they anticipate

requests and should generate routes on demand only in order to satisfy unanticipated route requests. Hence, a single route server can distribute its computational burden over time.

- Route servers should cache the results of route generation, in order to minimize the computation associated with responding to future route requests.
- To handle requested service limits, a route server should always select the first route generated that satisfies all of the requested service limits.
- To handle multi-criteria optimization in route selection, a route server should generate routes that are optimal with respect to the first optimal requested service listed in the ROUTE REQUEST message. The route server should resolve ties between otherwise equivalent routes by evaluating these routes according to the other optimal requested services contained in the ROUTE REQUEST message, in the order in which they are listed. With respect to the route server's routing information database, the selected route is optimal according to the first optimal requested service listed in the ROUTE REQUEST message but is not necessarily optimal according to any other optimal requested service listed in the ROUTE REQUEST message.

ti 2 - To handle a mixture of requested service limits and optimal requested services, a route server should generate routes that satisfy all of the requested service limits. The route server should resolve ties between otherwise equivalent routes by evaluating these routes as described in the multi-criteria optimization case above.

ti 2 - All else being equal, a route server should always prefer minimum-hop routes, because they minimize the amount of network resources consumed by the routes.

ti 2 - A route server should generate at least one route to each component of a partitioned destination domain, because it may not know in which domain component the destination host resides. Hence, a route server can maximize the chances of providing a feasible route to a destination within a partitioned domain.

6.1 Searching

All domains need not execute the identical route generation procedure. Each domain administrator is free to specify the IDPR route generation procedure for route servers in its own domain, making the procedure as simple or as complex as desired.

We offer an IDPR route generation procedure as a model. With slight modification, this procedure can be made to search in either BF or SPF order. The procedure can be used either to generate a single policy route from the source to a specified destination domain or to generate a set of policy routes from the source domain to all destination domains. If the source or destination domain has a proxy, then the source or destination endpoint of the policy route is a proxy domain and not the actual source or destination domain.

For high-bandwidth traffic flows, BF search is the recommended search technique, because it produces minimum-hop routes. For low-bandwidth traffic flows, the route server may use either BF search or SPF search. The computational complexity of BF search is $O(N + L)$ and hence it is the search procedure of choice, except when generating routes with optimal requested services. We recommend using SPF search only for optimal requested services and never in response to a request for a maximum bandwidth route.

6.1.1. Implementation

Data Structures:

The routing information database contains the graph of an internetwork, in which virtual gateways are the nodes and intra-domain routes between virtual gateways are the links. During route generation, each route is represented as a sequence of virtual gateways, domains, and relevant transit policies, together with a list of route characteristics, stored in a temporary array and indexed by destination domain.

- Execute the Policy Consistency routine, first with the source domain the given domain and second with the destination domain as the given domain. If any policy inconsistency precludes the requested traffic flow, go to Exit.
- For each domain, initialize a null route, set the route bandwidth to and set the following route characteristics to infinity: route delay, route delay variation, session monetary cost, and route length in hops.
- With each operational virtual gateway in the source or source proxy domain, associate the initial route characteristics.
- Initialize a next-node data structure which will contain, for each route in progress, the virtual gateway at the current endpoint of the route together with the associated route characteristics. The next-node data structure determines the order in which routes get expanded.

BF: A fifo queue.

SPF: A heap, ordered according to the first optimal requested service listed in the ROUTE REQUEST message.

Remove Next Node: These steps are performed for each virtual gateway in the next-node data structure.

- If there are no more virtual gateways in the next-node data structure, go to Exit.
- Extract a virtual gateway and its associated route characteristics from the next-node data structure, obtain the adjacent domain, and:

SPF: Remake the heap.

- If there is a specific destination domain and if for the primary optimal service:

BF: Route length in hops.

SPF: First optimal requested service listed in the ROUTE REQUEST message.

the extracted virtual gateway's associated route characteristic is no better than that of the destination domain, go to Remove Next Node.

- Execute the Policy Consistency routine with the adjacent domain as given domain. If any policy inconsistency precludes the requested traffic flow, go to Remove Next Node.
- Check that the source domain's transit policies do not preclude traffic generated by members of the source host set with the specified user class and requested services, from flowing to the adjacent domain as destination. This check is necessary because the route server caches what it considers to be all feasible routes, to intermediate destination domains, generated during the computation of the requested route. If there are no policy inconsistencies, associate the route and its characteristics with the adjacent domain as destination.
- If there is a specific destination domain and if the adjacent domain is the destination or destination proxy domain, go to Remove Next Node.
- Record the set of all exit virtual gateways in the adjacent

domain which the adjacent domain's transit policies permit the requested traffic flow and which are currently reachable from the entry virtual gateway.

Next Node:

These steps are performed for all exit virtual gateways in the above set.

- If there are no exit virtual gateways in the set, go to Remove Next Node.
- Compute the characteristics for the route to the exit virtual gateway, and check that all of the route characteristics are within the requested service limits. If any of the route characteristics are outside of these limits, go to Next Node.
- Compare these route characteristics with those already associated with the exit virtual gateway (there may be none, if this is the first time the exit virtual gateway has been visited in the search), according to the primary optimal service.
- Select the route with the optimal value of the primary optimal service, resolve ties by considering optimality according to any other optimal requested services in the order in which they are listed in the ROUTE REQUEST message, and associate the selected route and its characteristics with the exit virtual gateway.
- Add the virtual gateway to the next-node structure:
 - BF: Add to the end of the fifo queue.
 - SPF: Add to the heap.
 - and go to Next Node.

Exit:

Return a response to the route request, consisting of either a set of candidate policy routes or an indication that the route request cannot be fulfilled.

Policy Consistency: Check policy consistency for the given domain.

- Check that the given domain is not specified as an excluded domain in the route request.
- Check that the given domain's transit policies do not preclude traffic generated by members of the source host set with the

specified user class and requested services, from flowing to the destination domain.

During the computation of the requested routes, a route server also caches what it considers to be all feasible routes to intermediate destination domains, thus increasing the chances of being able to respond to a future route request without having to generate a new route. The route server does perform some policy consistency checks on the routes, as they are generated, to intermediate destinations. However, these routes may not in fact be feasible; the transit domains contained on the routes may not permit traffic between the source and the given intermediate destinations. Hence, before dispensing such a route in response to a route request, a route server must check that the transit policies of the constituent domains are consistent with the source and destination of the traffic flow.

6.2. Route Directionality

A path agent may wish to set up a bidirectional path using a route supplied by a route server. (Refer to sections 7.2 and 7.4 for detailed discussions of path directionality.) However, a route server can only guarantee that the routes it supplies are feasible if used in the direction from source to destination. The reason is that the route server, which resides in the source or source proxy domain, does not have access to, and thus cannot account for, the source policies of the destination domain. Nevertheless, the route server can provide the path agent with an indication of its assessment of route feasibility in the direction from destination to source.

A necessary but insufficient condition for a route to be feasible in the direction from destination to source is as follows. The route must be consistent, in the direction from destination to source, with the transit policies of the domains that compose the route. The transit policy consistency checks performed by the route server during route generation account for the direction from source to destination but not for the direction from destination to source. Only after a route server generates a feasible route from source to destination does it perform the transit policy consistency checks for the route in the direction from destination to source. Following these checks, the route server includes in its ROUTE RESPONSE message to the path agent an indication of its assessment of route feasibility in each direction.

6.3. Route Database

A policy route, as originally specified by a route server, is an ordered list of virtual gateways, domains, and transit policies: VG 1 - AD 1 - TP 1 - ... - VG n - AD n - TP n. where VG i is the virtual gateway that serves as exit from AD i-1 and entry to AD i, and TP i is the set of transit policies associated with AD i and relevant to the particular route. Each route is indexed by source and destination domain. Route servers and paths agents store policy routes in route databases maintained as caches whose entries must be periodically flushed to avoid retention of stale policy routes. A route server's route database is the set of all routes it has generated on behalf of its domain as source or source proxy; associated with each route in the database are its route characteristics. A path agent's route database is the set of all routes it has requested and received from route servers on behalf of hosts for which it is configured to act.

When attempting to locate a feasible route for a traffic flow, a path agent first consults its own route database before querying a route server. If the path agent's route database contains one or more routes between the given source and destination domains and accommodating the given host set and UCI, then the path agent checks each such route against the set of excluded domains listed in the source policy. The path agent either selects the first route encountered that does not include the excluded domains, or, if no such route exists in its route database, requests a route from a route server.

A path agent must query a route server for routes when it is unable to fulfill a route request from its own route database. Moreover, we recommend that a path agent automatically forward to a route server, all route requests with non-null requested services. The reason is that the path agent retains no route characteristics in its route database. Hence, the path agent cannot determine whether an entry in its route database satisfies the requested services.

When responding to a path agent's request for a policy route, a route server first consults its route database, unless the ROUTE REQUEST message contains an explicit directive to generate a new route. If its route database contains one or more routes between the given source and destination domains and accommodating the given host set and UCI, the route server checks each such route against the set of excluded domains listed in the ROUTE REQUEST message. The route server either selects all routes encountered that do not include the excluded domains, or, if no such route exists in its route database, attempts to generate such a route. Once the route server selects a set of routes, it then checks each such route against the services

requested by the path agent and the services offered by the domains composing the route. To obtain the offered services information, the route server consults its routing information database. The route server either selects the first route encountered that is consistent with both the requested and offered services, or, if no such route exists in its route database, attempts to generate such a route.

6.3.1. Cache Maintenance

Each route stored in a route database has a maximum cache lifetime equal to `rdb_rs` minutes for a route server and `rdb_ps` minutes for a path agent. Route servers and path agents reclaim cache space by flushing entries that have attained their maximum lifetimes. Moreover, paths agents reclaim cache space for routes whose paths have failed to be set up successfully or have been torn down (see section 7.4).

Nevertheless, cache space may become scarce, even with reclamation of entries. If a cache fills, the route server or path agent logs the event for network management. To obtain space in the cache when the cache is full, the route server or path agent deletes from the cache the oldest entry.

7. Path Control Protocol and Data Message Forwarding Procedure

Two entities in different domains may exchange IDPR data messages, only if there exists an IDPR path set up between the two domains. Path setup requires cooperation among path agents and intermediate policy gateways. Path agents locate policy routes, initiate the Path Control Protocol (PCP), and manage existing paths between administrative domains. Intermediate policy gateways verify that a given policy route is consistent with their domains' transit policies, establish the forwarding information, and forward messages along existing paths.

Each policy gateway and each route server contains a path agent. The path agent that initiates path setup in the source or source proxy domain is the "originator", and the path agent that handles the originator's path setup message in the destination or destination proxy domain is the "target". Every path has two possible directions of traffic flow: from originator to target and from target to originator. Path control messages are free to travel in either direction, but data messages may be restricted to only one direction.

Once a path for a policy route is set up, its physical realization is a set of consecutive policy gateways, with policy gateways or route servers forming the endpoints. Two successive entities in this set belong to either the same domain or the same virtual gateway. A

policy gateway or route server may, at any time, recover the resources dedicated to a path that goes through it by tearing down that path. For example, a policy gateway may decide to tear down a path that has not been used for some period of time.

PCP may build multiple paths between source and destination domains, but it is not responsible for managing such paths as a group or for eliminating redundant paths.

7.1. An Example of Path Setup

We illustrate how path setup works by stepping through an example. Suppose host Hx in domain AD X wants to communicate with host Hy in domain AD Y and that both AD X and AD Y support IDPR. Hx need not know the identity of its own domain or of Hy's domain in order to send messages to Hy. Instead, Hx simply forwards a message bound for Hy to one of the gateways on its local network, according to its local forwarding information only. If the recipient gateway is a policy gateway, the resident path agent determines how to forward the message outside of the domain. Otherwise, the recipient gateway forwards the message to another gateway in AD X, according to its local forwarding information. Eventually, the message will arrive at a policy gateway in AD X, as policy gateways are the only egress points to other domains, in domains that support IDPR.

The path agent resident in the recipient policy gateway uses the message header, including source and destination addresses and any requested service information (for example, type of service), in order to determine whether it is an intra-domain or inter-domain message, and if inter-domain, whether it requires an IDPR policy route. Specifically, the path agent attempts to locate a forwarding information database entry for the given traffic flow, from the information contained in the message header. In the future, for IP messages, the relevant header information may also include special service-specific IP options or even information from higher layer protocols.

Forwarding database entries exist for all of the following:

- All intra-domain traffic flows. Intra-domain forwarding information is integrated into the forwarding information database as soon as it is received.
- Inter-domain traffic flows that do not require IDPR policy routes. Non-IDPR forwarding information is integrated into the forwarding database as soon as it is received.
- IDPR inter-domain traffic flows for which a path has already been

set up. IDPR forwarding information is integrated into the forwarding database only during path setup.

The path agent uses the message header contents to guide the search for a forwarding information database entry for a given traffic flow. We recommend a radix search to locate such an entry. When the search terminates, it produces either an entry, or, in the case of a new IDPR traffic flow, a directive to generate an entry. If the search terminates in an existing forwarding information database entry, the path agent forwards the message according to that entry.

Suppose that the search terminates indicating that the traffic flow from Hx to Hy requires an IDPR policy route and that no entry in the forwarding information database yet exists for that traffic flow. In this case, the path agent first determines the source and destination domains associated with the message's source and destination addresses, before attempting to obtain a policy route. The path agent relies on the mapping servers to supply the domain information, but it caches all mapping server responses locally to limit the number of future queries. When attempting to resolve an address to a domain, the path agent always checks its local cache before contacting a mapping server.

After obtaining the domain information, the path agent attempts to obtain a policy route to carry the traffic from Hx to Hy. The path agent relies on route servers to supply policy routes, but it caches all route server responses locally to limit the number of future queries. When attempting to locate a suitable policy route, the path agent usually consults its local cache before contacting a route server, as described previously in section 6.3.

If no suitable cache entry exists, the path agent queries the route server, providing it with the source and destination domains together with source policy information carried in the host message or specified through configuration. Upon receiving a policy route query, a route server consults its route database. If it cannot locate a suitable route in its route database, the route server attempts to generate at least one route to AD Y, consistent with the requested services for Hx.

The route server always returns a response to the path agent, regardless of whether it is successful in locating a suitable policy route. The response to a successful route query consists of a set of candidate routes, from which the path agent makes its selection. We expect that a path agent will normally choose a single route from a candidate set. Nevertheless, IDPR does not preclude a path agent from selecting multiple routes from the candidate set. A path agent may desire multiple routes to support features such as fault

tolerance or load balancing; however, IDPR does not currently specify how the path agent should use multiple routes.

If the policy route is a new route provided by the route server, there will be no existing path for the route, and thus the path agent must set up such a path. However, if the policy route is an existing route extracted from the path agent's cache, there may well be an existing path for the route, set up to accommodate a host traffic flow. IDPR permits multiple traffic flows to use the same path, provided that all traffic flows sharing the path travel between the same endpoint domains and have the same service requirements. Nevertheless, IDPR does not preclude a path agent from setting up distinct paths along the same policy route to preserve the distinction between host traffic flows.

The path agent associates an identifier with the path, which is included in each message that travels down the path and is used by the policy gateways along the path in order to determine how to forward the message. If the path already exists, the path agent uses the preexisting identifier. However, for new paths, the path agent chooses a path identifier that is different from those of all other paths that it manages. The path agent also updates its forwarding information database to reference the path identifier and modifies its search procedure to yield the correct entry in the forwarding information database given the data message header.

For new paths, the path agent initiates path setup, communicating the policy route, in terms of requested services, constituent domains, relevant transit policies, and the connecting virtual gateways, to policy gateways in intermediate domains. Using this information, an intermediate policy gateway determines whether to accept or refuse the path and to which next policy gateway to forward the path setup information. The path setup procedure allows policy gateways to set up a path in both directions simultaneously. Each intermediate policy gateway, after path acceptance, updates its forwarding information database to include an entry that associates the path identifier with the appropriate previous and next hop policy gateways.

When a policy gateway in AD Y accepts a path, it notifies the source path agent in AD X. We expect that the source path agent will normally wait until a path has been successfully established before using it to transport data traffic. However, PCP does not preclude a path agent from forwarding messages along a path prior to confirmation of successful path establishment. Paths remain in place until they are torn down because of failure, expiration, or when resources are scarce, preemption in favor of other paths.

We note that data communication between Hx and Hy may occur over two separate IDPR paths: one from AD X to AD Y and one from AD Y to AD X. The reasons are that within a domain, hosts know nothing about path agents nor IDPR paths, and path agents know nothing about other path agents' existing IDPR paths. Thus, in AD Y, the path agent that terminates the path from AD X may not be the same as the path agent that receives traffic from Hy destined for Hx. In this case, receipt of traffic from Hy forces the second path agent to set up an independent path from AD Y to AD X.

7.2. Path Identifiers

Each path has an associated path identifier, unique throughout an internetwork. Every IDPR data message travelling along that path includes the path identifier, used for message forwarding. The path identifier is the concatenation of three items: the identifier of the originator's domain, the identifier of the originator's policy gateway or route server, and a 32-bit local path identifier specified by the originator. The path identifier and the CMTTP transaction identifier have analogous syntax and play analogous roles in their respective protocols.

When issuing a new path identifier, the originator always assigns a local path identifier that is different from that of any other active or recently torn-down path originally set up by that path agent. This helps to distinguish new paths from replays. Hence, the originator must keep a record of each extinct path for long enough that all policy gateways on the path will have eliminated any reference to it from their memories. The right-most 30 bits of the local identifier are the same for each path direction, as they are assigned by the originator. The left-most 2 bits of the local identifier indicate the path direction.

At path setup time, the originator specifies which of the path directions to enable contingent upon the information received from the route server in the ROUTE RESPONSE message. By "enable", we mean that each path agent and each intermediate policy gateway establishes an association between the path identifier and the previous and next policy gateways on the path, which it uses for forwarding data messages along that path. IDPR data messages may travel in the enabled path directions only, but path control messages are always free to travel in either path direction. The originator may enable neither path direction, if the entire data transaction can be carried in the path setup message itself. In this case, the path agents and the intermediate policy gateways do not establish forwarding associations for the path, but they do verify consistency of the policy information contained in the path setup message, with their own transit policies, before forwarding the setup message on to the

next policy gateway.

The path direction portion of the local path identifier has different interpretations, depending upon message type. In an IDPR path setup message, the path direction indicates the directions in which the path should be enabled: the value 01 denotes originator to target, the value 10 denotes target to originator, the value 11 denotes both directions, and the value 00 denotes neither direction. Each policy gateway along the path interprets the path direction in the setup message and sets up the forwarding information as directed. In an IDPR data message, the path direction indicates the current direction of traffic flow: either 01 for originator to target or 10 for target to originator. Thus, if for example, an originator sets up a path enabling only the direction from target to originator, the target sends data messages containing the path identifier selected by the originator together with the path direction set equal to 10.

Instead of using path identifiers that are unique throughout an internetwork, we could have used path identifiers that are unique only between a pair of consecutive policy gateways and that change from one policy gateway pair to the next. The advantage of locally unique path identifiers is that they may be much shorter than globally unique identifiers and hence consume less transmission bandwidth. However, the disadvantage is that the path identifier carried in each IDPR data message must be modified at each policy gateway, and hence if the integrity/authentication information covers the path identifier, it must be recomputed at each policy gateway. For security reasons, we have chosen to include the path identifier in the set of information covered by the integrity/authentication value, and moreover, we advocate public-key based signatures for authentication. Thus, it is not possible for intermediate policy gateways to modify the path identifier and then recompute the correct integrity/authentication value. Therefore, we have decided in favor of path identifiers that do not change from hop to hop and hence must be globally unique. To speed forwarding of IDPR data messages with long path identifiers, policy gateways should hash the path identifiers in order to index IDPR forwarding information.

7.3. Path Control Messages

Messages exchanged by the path control protocol are classified into "requests": SETUP, TEARDOWN, REPAIR; and "responses": ACCEPT, REFUSE, ERROR. These messages have significance for intermediate policy gateways as well as for path agents.

SETUP:

Establishes a path by linking together pairs of policy gateways. The SETUP message is generated by the originator and propagates

to the target. In response to a SETUP message, the originator expects to receive an ACCEPT, REFUSE, or ERROR message. The SETUP message carries all information necessary to set up the path including path identifier, requested services, transit policy information relating to each domain traversed, and optionally, expedited data.

ACCEPT: Signals successful path establishment. The ACCEPT message is generated by the target, in response to a SETUP message, and propagates back to the originator. Reception of an ACCEPT message by the originator indicates that the originator can now safely proceed to send data along the path. The ACCEPT message contains the path identifier and an optional reason for conditional acceptance.

REFUSE: Signals that the path could not be successfully established, either because resources were not available or because there was an inconsistency between the services requested by the source and the services offered by a transit domain along the path. The REFUSE message is generated by the target or by any intermediate policy gateway, in response to a SETUP message, and propagates back to the originator. All recipients of a REFUSE message recover the resources dedicated to the given path. The REFUSE message contains the path identifier and the reason for path refusal.

TEARDOWN: Tears down a path, typically when a non-recoverable failure is detected. The TEARDOWN message may be generated by any path agent or policy gateway in the path and usually propagates in both path directions. All recipients of a TEARDOWN message recover the resources dedicated to the given path. The TEARDOWN message contains the path identifier and the reason for path teardown.

REPAIR: Establishes a repaired path by linking together pairs of policy gateways. The REPAIR message is generated by a policy gateway after detecting that the next policy gateway on one of its existing paths is unreachable. A policy gateway that generates a REPAIR message propagates the message forward at most one virtual gateway. In response to a REPAIR message, the policy gateway expects to receive an ACCEPT, REFUSE, TEARDOWN, or ERROR message. The REPAIR message carries the original SETUP message.

ERROR: Transports information about a path error back to the originator, when a PCP message contains unrecognized information. The ERROR message may be generated by the target or by any intermediate policy gateway and propagates back to the

originator. Most, but not all, ERROR messages are generated in response to errors encountered during path setup. The ERROR message includes the path identifier and an explanation of the error detected.

Policy gateways use CMTP for reliable transport of PCP messages, between path agents and policy gateways and between consecutive policy gateways on a path. PCP must communicate to CMTP the maximum number of transmissions per path control message, `pcp_ret`, and the interval between path control message retransmissions, `pcp_int` microseconds. All path control messages, except ERROR messages, may be transmitted up to `pcp_ret` times; ERROR messages are never retransmitted. A path control message is "acceptable" if:

- It passes the CMTP validation checks.
- Its timestamp is less than `pcp_old` (300) seconds behind the recipient's internal clock time.
- It carries a recognized path identifier, provided it is not a SETUP message.

An intermediate policy gateway on a path forwards acceptable PCP messages. As we describe in section 7.4 below, SETUP messages must undergo additional tests at each intermediate policy gateway prior to forwarding. Moreover, receipt of an acceptable ACCEPT, REFUSE, TEARDOWN, or ERROR message at either path agent or at any intermediate policy gateway indirectly cancels any active local CMTP retransmissions of the original SETUP message. When a path agent or intermediate policy gateway receives an unacceptable path control message, it discards the message and logs the event for network management. The path control message age limit reduces the likelihood of denial of service attacks based on message replay.

7.4. Setting Up and Tearing Down a Path

Path setup begins when the originator generates a SETUP message containing:

- The path identifier, including path directions to enable.
- An indication of whether the message includes expedited data.
- The source user class identifier.
- The requested services (see section 5.5.2) and source-specific information (see section 7.6.1) for the path.

- For each domain on the path, the domain component, applicable transit policies, and entry and exit virtual gateways.

The only mandatory requested service is the maximum path lifetime, `pth_lif`, and the only mandatory source-specific information is the data message integrity/authentication type. If these are not specified in the path setup message, each recipient policy gateway assigns them default values, (60) minutes for `pth_lif` and no authentication for integrity/authentication type. Each path agent and intermediate policy gateway tears down a path when the path lifetime is exceeded. Hence, no single source can indefinitely monopolize policy gateway resources or still functioning parts of partially broken paths.

After generating the SETUP message and establishing the proper local forwarding information, the originator selects the next policy gateway on the path and forwards the SETUP message to the selected policy gateway. The next policy gateway selection procedure, described below, applies when either the originator or an intermediate policy gateway is making the selection. We have elected to describe the procedure from the perspective of a selecting intermediate policy gateway.

The policy gateway selects the next policy gateway on a path, in round-robin order from its list of policy gateways contained in the present or next virtual gateway, as explained below. In selecting the next policy gateway, the policy gateway uses information contained in the SETUP message and information provided by VGP and by the intra-domain routing procedure.

If the selecting policy gateway is a domain entry point, the next policy gateway must be:

- A member of the next virtual gateway listed in the SETUP message.
- Reachable according to intra-domain routes supporting the transit policies listed in the SETUP message.
- Able to reach, according to VGP, the next domain component listed in the SETUP message.

In addition, the selecting policy gateway may use quality of service information supplied by intra-domain routing to resolve ties between otherwise equivalent next policy gateways in the same domain. In particular, the selecting policy gateway may select the next policy gateway whose connecting intra-domain route is optimal according to the requested services listed in the SETUP message.

If the selecting policy gateway is a domain exit point, the next policy gateway must be:

- A member of the current virtual gateway listed in the SETUP message (which is also the selecting policy gateway's virtual gateway).
- Reachable according to VGP.
- A member of the next domain component listed in the SETUP message.

Once the originator or intermediate policy gateway selects a next policy gateway, it forwards the SETUP message to the selected policy gateway. Each recipient (policy gateway or target) of an acceptable SETUP message performs several checks on the contents of the message, in order to determine whether to establish or reject the path. We describe these checks in detail below from the perspective of a policy gateway as SETUP message recipient.

7.4.1. Validating Path Identifiers

The recipient of a SETUP message first checks the path identifier, to make sure that it does not correspond to that of an already existing or recently extinct path. To detect replays, malicious or otherwise, path agents and policy gateways maintain a record of each path that they establish, for $\max\{\text{pth_lif}, \text{pcp_old}\}$ seconds. If the path identifier and timestamp carried in the SETUP message match a stored path identifier and timestamp, the policy gateway considers the message to be a retransmission and does not forward the message. If the path identifier carried in the SETUP message matches a stored path identifier but the two timestamps do not agree, the policy gateway abandons path setup, logs the event for network management, and returns an ERROR message to the originator via the previous policy gateway.

7.4.2. Path Consistency with Configured Transit Policies

Provided the path identifier in the SETUP message appears to be new, the policy gateway proceeds to determine whether the information contained within the SETUP message is consistent with the transit policies configured for its domain. The policy gateway must locate the source and destination domains, the source host set and user class identifier, and the domain-specific information for its own domain, within the SETUP message, in order to evaluate path consistency. If the policy gateway fails to recognize the source user class (or one or more of the requested services), it logs the event for network management but continues with path setup. If the policy gateway fails to locate its domain within the SETUP message, it abandons path setup, logs the event for network management, and

returns an ERROR message to the originator via the previous policy gateway. The originator responds by tearing down the path and subsequently removing the route from its cache.

Once the policy gateway locates its domain-specific portion of the SETUP message, it may encounter the following problems with the contents:

- The domain-specific portion lists a transit policy not configured for the domain.
- The domain-specific portion lists a virtual gateway not configured for the domain.

In each case, the policy gateway abandons path setup, logs the event for network management, and returns an ERROR message to the originator via the previous policy gateway. These types of ERROR messages indicate to the originator that the route may have been generated using information from an out-of-date CONFIGURATION message.

The originator reacts to the receipt of such an ERROR message as follows. First, it tears down the path and removes the route from its cache. Then, it issues to a route server a ROUTE REQUEST message containing a directive to refresh the routing information database, with the most recent CONFIGURATION message from the domain that issued the ERROR message, before generating a new route.

Once it verifies that its domain-specific information in the SETUP message is recognizable, the policy gateway then checks that the information contained within the SETUP message is consistent with the transit policies configured for its domain. A policy gateway at the entry to a domain checks path consistency in the direction from originator to target, if the enabled path directions include originator to target. A policy gateway at the exit to a domain checks path consistency in the direction from target to originator, if the enabled path directions include target to originator.

When evaluating the consistency of the path with the transit policies configured for the domain, the policy gateway may encounter any of the following problems with SETUP message contents:

- A transit policy does not apply in the given direction between the virtual gateways listed in the SETUP message.
- A transit policy denies access to traffic from the given host set between the source and destination domains listed in the SETUP message.

- A transit policy denies access to traffic from the source user class listed in the SETUP message.
- A transit policy denies access to traffic at the current time.

In each case, the policy gateway abandons path setup, logs the event for network management, and returns a REFUSE message to the originator via the previous policy gateway. These types of REFUSE messages indicate to the originator that the route may have been generated using information from an out-of-date CONFIGURATION message. The REFUSE message also serves to teardown the path.

The originator reacts to the receipt of such a REFUSE message as follows. First, it removes the route from its cache. Then, it issues to a route server a ROUTE REQUEST message containing a directive to refresh the routing information database, with the most recent CONFIGURATION message from the domain that issued the REFUSE message, before generating a new route.

7.4.3. Path Consistency with Virtual Gateway Reachability

Provided the information contained in the SETUP message is consistent with the transit policies configured for its domain, the policy gateway proceeds to determine whether the path is consistent with the reachability of the virtual gateway containing the potential next hop. To determine virtual gateway reachability, the policy gateway uses information provided by VGP and by the intra-domain routing procedure.

When evaluating the consistency of the path with virtual gateway reachability, the policy gateway may encounter any of the following problems:

- The virtual gateway containing the potential next hop is down.
- The virtual gateway containing the potential next hop is not reachable via any intra-domain routes supporting the transit policies listed in the SETUP message.
- The next domain component listed in the SETUP message is not reachable.

Each of these determinations is made from the perspective of a single policy gateway and may not reflect actual reachability. In each case, the policy gateway encountering such a problem returns a REFUSE message to the previous policy gateway which then selects a different next policy gateway, in round-robin order, as described in previously. If the policy gateway receives the same response from

all next policy gateways selected, it abandons path setup, logs the event for network management, and returns the REFUSE message to the originator via the previous policy gateway. These types of REFUSE messages indicate to the originator that the route may have been generated using information from an out-of-date DYNAMIC message. The REFUSE message also serves to teardown the path.

The originator reacts to the receipt of such a REFUSE message as follows. First, it removes the route from its cache. Then, it issues to a route server a ROUTE REQUEST message containing a directive to refresh the routing information database, with the most recent DYNAMIC message from the domain that issued the REFUSE message, before generating a new route.

7.4.4. Obtaining Resources

Once the policy gateway determines that the SETUP message contents are consistent with the transit policies and virtual gateway reachability of its domain, it attempts to gain resources for the new path. For this version of IDPR, path resources consist of memory in the local forwarding information database. However, in the future, path resources may also include reserved link bandwidth.

If the policy gateway does not have sufficient resources to establish the new path, it uses the following algorithm to determine whether to generate a REFUSE message for the new path or a TEARDOWN message for an existing path in favor of the new path. There are two cases:

- No paths have been idle for more than `pcp_idle` (300) seconds. In this case, the policy gateway returns a REFUSE message to the previous policy gateway. This policy gateway then tries to select a different next policy gateway, as described previously, provided the policy gateway that issued the REFUSE message was not the target. If the REFUSE message was issued by the target or if there is no available next policy gateway, the policy gateway returns the REFUSE message to the originator via the previous policy gateway and logs the event for network management. The REFUSE message serves to tear down the path.
- At least one path has been idle for more than `pcp_idle` seconds. In this case, the policy gateway tears down an older path in order to accommodate the newer path and logs the event for network management. Specifically, the policy gateway tears down the least recently used path among those that have been idle for longer than `pcp_idle` seconds, resolving ties by choosing the oldest such path.

If the policy gateway has sufficient resources to establish the path,

it attempts to update its local forwarding information database with information about the path identifier, previous and next policy gateways on the path, and directions in which the path should be enabled for data traffic transport.

7.4.5 Target Response

When an acceptable SETUP message successfully reaches an entry policy gateway in the destination or destination proxy domain, this policy gateway performs all of the SETUP message checks described in the above sections. The policy gateway's path agent then becomes the target, provided no checks fail, unless there is an explicit target specified in the SETUP message. For example, remote route servers act as originator and target during RSQP message exchanges (see section 5.2). If the recipient policy gateway is not the target, it attempts to forward the SETUP message to the target along an intra-domain route. However, if the target is not reachable via intra-domain routing, the policy gateway abandons path setup, logs the event for network management, and returns a REFUSE message to the originator via the previous policy gateway. The REFUSE message serves to tear down the path.

Once the SETUP message reaches the target, the target determines whether it has sufficient path resources. The target generates an ACCEPT message, provided it has sufficient resources to establish the path. Otherwise, it generates a REFUSE message.

The target may choose to use the reverse path to transport data traffic to the source domain, if the enabled path directions include 10 or 11. However, the target must first verify the consistency of the reverse path with its own domain's configured transit policies, before sending data traffic over that path.

7.4.6. Originator Response

The originator expects to receive an ACCEPT, REFUSE, or ERROR message in response to a SETUP message and reacts as follows:

- The originator receives an ACCEPT message, confirming successful path establishment. To expedite data delivery, the originator may forward data messages along the path prior to receiving an ACCEPT message, with the understanding that there is no guarantee that the path actually exists.
- The originator receives a REFUSE message or an ERROR message, implying that the path could not be successfully established. In response, the originator attempts to set up a different path to the same destination, as long as the number of selected different paths

does not exceed `setup_try` (3). If the originator is unsuccessful after `setup_try` attempts, it abandons path setup and logs the event for network management.

- The originator fails to receive any response to the SETUP message within `setup_int` microseconds after transmission. In this case, the originator attempts path setup using the same policy route and a new path identifier, as long as the number of path setup attempts using the same route does not exceed `setup_ret` (2). If the originator fails to receive a response to a SETUP message after `setup_ret` attempts, it logs the event for network management and then proceeds as though it received a negative response, namely a REFUSE or an ERROR, to the SETUP message. Specifically, it attempts to set up a different path to the same destination, or it abandons path setup altogether, depending on the value of `setup_try`.

7.4.7. Path Life

Once set up, a path does not live forever. A path agent or policy gateway may tear down an existing path, provided any of the following conditions are true:

- The maximum path lifetime (in minutes, bytes, or messages) has been exceeded at the originator, the target, or an intermediate policy gateway. In each case, the IDPR entity detecting path expiration logs the event for network management and generates a TEARDOWN message as follows:
 - o The originator path agent generates a TEARDOWN message for propagation toward the target.
 - o The target path agent generates a TEARDOWN message for propagation toward the originator.
 - o An intermediate policy gateway generates two TEARDOWN messages, one for propagation toward the originator and one for propagation toward the target.
- The previous or next policy gateway becomes unreachable, across a virtual gateway or across a domain according to a given transit policy, and the path is not reparable. In either case, the policy gateway detecting the reachability problem logs the event for network management and generates a TEARDOWN message as follows:
 - o If the previous policy gateway is unreachable, an intermediate policy gateway generates a TEARDOWN message for propagation to the target.

- o If the next policy gateway is unreachable, an intermediate policy gateway generates a TEARDOWN message for propagation to the originator.
- All of the policy gateway's path resources are in use at the originator, the target, or an intermediate policy gateway, a new path requires resources, and the given existing path is expendable, according to the least recently used criterion discussed in section 7.4.4 above. In each case, the IDPR entity initiating path preemption logs the event for network management and generates a TEARDOWN message as follows:
 - o The originator path agent generates a TEARDOWN message for propagation toward the originator.
 - o The target path agent generates a TEARDOWN message for propagation toward the originator.
 - o An intermediate policy gateway generates two TEARDOWN messages, one for propagation toward the originator and one for propagation toward the target.

Path teardown at a path agent or policy gateway, whether initiated by one of the above events, by receipt of a TEARDOWN message, or by receipt of a REFUSE message during path setup (as discussed in the previous sections), results in the path agent or policy gateway releasing all resources devoted to both directions of the path.

7.5. Path Failure and Recovery

When a policy gateway fails, it may not be able to save information pertaining to its established paths. Thus, when the policy gateway returns to service, it may have no recollection of the paths set up through it and hence may no longer be able to forward data messages along these paths. We expect that when a policy gateway fails, it will usually be out of service for long enough that the up/down protocol and the intra-domain routing procedure can detect that the particular policy gateway is no longer reachable. In this case, adjacent or neighbor policy gateways that have set up paths through the failed policy gateway and that have detected the failure, attempt local path repair (see section 7.5.2 below), and if unsuccessful, issue TEARDOWN messages for all affected paths.

7.5.1. Handling Implicit Path Failures

Nevertheless, policy gateways along a path must be able to handle the case in which a policy gateway fails and subsequently returns to service without either the up/down protocol or the intra-domain routing procedure detecting the failure; we do not expect this event to occur often. If the policy gateway, prior to failure, contained forwarding information for several established paths, it may now receive many IDPR data messages containing unrecognized path identifiers. The policy gateway should alert the data sources that their paths through it are no longer viable.

Policy gateways that receive IDPR data messages with unrecognized path identifiers take one of the following two actions, depending upon their past failure record:

- The policy gateway has not failed in the past `pg_up` (24) hour period. In this case, there are at least four possible reasons for the unrecognized path identifier in the data message:
 - o The data message path identifier has been corrupted in a way that is not detectable by the integrity/authentication value, if one is present.
 - o The policy gateway has experienced a memory error.
 - o The policy gateway failed sometime during the life of the path and source sent no data on the path for a period of `pg_up` hours following the failure. Although paths may persist for more than `pg_up` hours, we expect that they will also be used more frequently than once every `pg_up` hours.
 - o The path was not successfully established, and the originator sent data messages down the path prior to receiving a response to its SETUP message.

In all cases, the policy gateway discards the data message and logs the event for network management.

- The policy gateway has failed at least once in the past `pg_up` hour period. Thus, the policy gateway assumes that the unrecognized path identifier in the data message may be attributed to its failure. In response to the data message, the policy gateway generates an ERROR message containing the unrecognized path identifier. The policy gateway then sends the ERROR message back to the entity from which it received the data message, which should be equivalent to the previous policy gateway on the path.

When the previous policy gateway receives such an ERROR message, it decides whether the message is acceptable. If the policy gateway does not recognize the path identifier contained in the ERROR message, it does not find the ERROR message acceptable and subsequently discards the message. However, if the policy gateway does find the ERROR message acceptable, it then determines whether it has already received an ACCEPT message for the given path. If the policy gateway has not received an ACCEPT message for that path, it discards the ERROR message and takes no further action.

If the policy gateway has received an ACCEPT message for that path, it then attempts path repair, as described in section 7.5.2 below. Only if path repair is unsuccessful does the previous policy gateway generate a TEARDOWN message for the path and return it to the originator. The TEARDOWN message includes the domain and virtual gateway containing the policy gateway that failed, which aids the originator in selecting a new path that does not include the domain containing the failed policy gateway. This mechanism ensures that path agents quickly discover and recover from disrupted paths, while guarding against unwarranted path teardown.

7.5.2. Local Path Repair

Failure of one of more entities on a given path may render the path unusable. If the failure is within a domain, IDPR relies on the intra-domain routing procedure to find an alternate route across the domain, which leaves the path unaffected. If the failure is in a virtual gateway, policy gateways must bear the responsibility of repairing the path. Policy gateways nearest to the failure are the first to recognize its existence and hence can react most quickly to repair the path.

Relinquishing control over path repair to policy gateways in other domains may be unacceptable to some domain administrators. The reason is that these policy gateways cannot guarantee construction of a path that satisfies the source policies of the source domain, as they have no knowledge of other domains' source policies.

Nevertheless, limited local path repair is feasible, without distributing either source policy information throughout an internetwork or detailed path information among policy gateways in the same domain or in the same virtual gateway. We say that a path is "locally reparable" if there exists an alternate route between two policy gateways, separated by at most one virtual gateway, on the path. This definition covers path repair in the presence of failed routes between consecutive policy gateways as well as failed policy gateways themselves.

An IDPR entity attempts local repair of an established path, in the direction from originator to target, immediately after detecting that the next policy gateway on the path is no longer reachable. To prevent multiple path repairs in response to the same failure, we have stipulated that path repair can only be initiated in the direction from originator to target. The IDPR entity initiating local path repair attempts to find an alternate path to the policy gateway immediately following the unreachable policy gateway on the path.

Local path repair minimizes the disruption of data traffic flow caused by certain types of failures along an established path. Specifically, local path repair can accommodate an individual failed policy gateway or failed direct connection between two adjacent policy gateways. However, it can only be attempted through virtual gateways containing multiple peer policy gateways. Local path repair is not designed to repair paths traversing failed virtual gateways or domain partitions. Whenever local path repair is impossible, the failing path must be torn down.

7.5.3. Repairing a Path

When an IDPR entity detects through an ERROR message that the next policy gateway has no knowledge of a given path, it generates a REPAIR message and forwards it to the next policy gateway. This REPAIR message will reestablish the path through the next policy gateway.

When an entity detects that the next policy gateway on a path is no longer reachable, it takes one of the following actions, depending upon whether the entity is a member of the next policy gateway's virtual gateway.

- If the entity is not a member of the next policy gateway's virtual gateway, then one of the following two conditions must be true:
 - o The next policy gateway has a peer that is reachable via an intra-domain route consistent with the requested services. In this case, the entity generates a REPAIR message containing the original SETUP message and forwards it to the next policy gateway's peer.
 - o The next policy gateway has no peers that are reachable via intra-domain routes consistent with the requested services. In this case, the entity tears down the path back to the originator.
- If the entity is a member of the next policy gateway's virtual

gateway, then one of the following four conditions must be true:

- o The next policy gateway has a peer that belongs to the same domain component and is directly-connected to and reachable from the entity. In this case, the entity generates a REPAIR message and forwards it to the next policy gateway's peer.
- o The next policy gateway has a peer that belongs to the same domain component, is not directly-connected to the entity, but is directly-connected to and reachable from one of the entity's peers, which in turn is reachable from the entity via an intra-domain route consistent with the requested services. In this case, the entity generates a REPAIR message and forwards it to its peer.
- o The next policy gateway has no operational peers within its domain component, but is directly-connected to and reachable from one of the entity's peers, which in turn is reachable from the entity via an intra-domain route consistent with the requested services. In this case, the entity generates a REPAIR message and forwards it to its peer.
- o The next policy gateway has no operational peers within its domain component, and the entity has no operational peers which are both reachable via intra-domain routes consistent with the requested services and directly-connected to and reachable from the next policy gateway. In this case, the entity tears down the path back to the originator.

A recipient of a REPAIR message takes the following steps, depending upon its relationship to the entity that issued the REPAIR message.

- The recipient and the issuing entity are in the same domain or in same virtual gateway. In this case, the recipient extracts the SETUP message contained within the REPAIR message and treats the message as it would any other SETUP message. Specifically, the recipient checks consistency of the path with its domain's transit policies and virtual gateway reachability. If there are unrecognized portions of the SETUP message, the recipient generates an ERROR message, and if there are path inconsistencies, the recipient generates a REFUSE message. In either case, the recipient returns the corresponding message to the policy gateway from which it received the REPAIR message. Otherwise, if the recipient accepts the REPAIR message, it updates its local forwarding information database accordingly and forwards the REPAIR message to a potential next policy gateway, according to the information contained in the enclosed SETUP message.

- The recipient and the issuing entity are in different domains and different virtual gateways. In this case, the recipient extracts the SETUP message from the REPAIR message and determines whether the associated path matches any of its established paths. If the path does not match an established path, the recipient generates a REFUSE message and returns it to the policy gateway from which it received the REPAIR message. In response to the receipt of a REFUSE message, the policy gateway tries a different next policy gateway.

The path is reparable, if a path match is discovered. In this case, the recipient updates the path entry in the local forwarding information database and issues an ACCEPT message to the policy gateway from which it received the REPAIR message, which in turn returns the message to the entity that issued the REPAIR message. The path is irreparable if all potential next policy gateways have been exhausted and a path match has yet to be discovered. In this case, the policy gateway that fails to locate a next policy gateway issues a TEARDOWN message to return to the originator.

An IDPR entity expects to receive an ACCEPT, TEARDOWN, REFUSE, or ERROR message in response to a REPAIR message and reacts to these responses differently as follows:

- The entity always returns a TEARDOWN message to the originator via previous policy gateway.
- The entity does not return an ACCEPT message to the originator, but receipt of such a message indicates that the path has been successfully repaired.
- The entity infers that the path is irreparable and subsequently tears down the path and logs the event for network management, upon receipt of a REFUSE or ERROR message or when no response to the REPAIR message arrives within `setup_int` microseconds.


When an entity detects that the previous policy gateway on a path becomes unreachable, it expects to receive a REPAIR message within `setup_wait` microseconds. If the entity does not receive a REPAIR message for the path within that time, it infers that the path is irreparable and subsequently tears down the path and logs the event for network management.

7.6. Path Control Message Formats

The path control protocol number is equal to 3. We describe the contents of each type of PCP message below.

7.6.1. SETUP

The SETUP message type is equal to 0.



PATH ID

(64 bits) Path identifier consisting of the numeric identifier for the originator's domain (16 bits), the numeric identifier for the originator policy gateway or route server (16 bits), the path direction (2 bits), and the local path identifier (30 bits).

SRC AD (16 bits) Numeric identifier for the source domain, which may be different from the originator domain if the originator domain is a proxy for the source.

HST SET (16 bits) Numeric identifier for the source's host set.

UCI (8 bits) Numeric identifier for the source user class. The value 0 indicates that there is no particular source user class.

UNUSED (8 bits) Not currently used; must be set equal to 0.

NUM RQS (16 bits) Number of requested services.

DST AD (16 bits) Numeric identifier for the destination domain, which may be different from the target domain if the target domain is a proxy for the destination.

TGT ENT (16 bits) Numeric identifier for the target entity. A value of 0 indicates that there is no specific target entity.

AD PTR (16 bits) Byte offset from the beginning of the message indicating the location of the beginning of the domain-specific information, contained in the right-most 15 bits. The left-most bit indicates whether the message includes expedited data (1 expedited data, 0 no expedited data).

RQS TYP (16 bits) Numeric identifier for a type of requested service or source-specific information. Valid requested services are described in section 5.5.2. Valid source source-specific information includes the following types:

12. MD4/RSA data message authentication (see [16]).

13. MD5/RSA data message authentication (see [17]).

14. Billing address (variable).

15. Charge number (variable).

RQS LEN (16 bits) Length of the requested service or source-specific information, in bytes, beginning with the next field.

RQS SRV (variable) Description of the requested service or source-specific information.

AD LEN (8 bits) Length of the information associated with a particular domain on the route, in bytes, beginning with the next field.

VG (8 bits) Numeric identifier for an exit virtual gateway.

ADJ AD (16 bits) Numeric identifier for an adjacent domain.


ADJ CMP (16 bits) Numeric identifier for a component of the adjacent domain. Used to aid a policy gateway in routing across a virtual gateway connected to a partitioned domain.

NUM TP (16 bits) Number of transit policies that apply to the section of the path traversing the given domain component.

TP (16 bits) Numeric identifier for a transit policy.

7.6.2. ACCEPT

The ACCEPT message type is equal to 1.



PATH ID


(64 bits) Path identifier contained in the original SETUP message.

RSN TYP (optional, 8 bits) Numeric identifier for the reason for conditional path acceptance.

REASON (optional, variable) Description of the reason for conditional path acceptance. Currently, no reasons have been defined.

7.6.3 REFUSE

The REFUSE message type is equal to 2.



PATH ID

(64 bits) Path identifier contained in the original SETUP message.

RSN TYP (8 bits) Numeric identifier for the reason for path refusal.


REASON (variable) Description of the reason for path refusal. Valid

reasons include the following types:

1. Transit policy does not apply between the virtual gateways in a given direction. Numeric identifier for the transit policy (16 bits).
2. Transit policy denies access to traffic from the host set between the source and destination domains. Numeric identifier for the transit policy (16 bits).
3. Transit policy denies access to traffic from the source user class. Numeric identifier for the transit policy (16 bits).
4. Transit policy denies access to traffic at the current time. Numeric identifier for the transit policy (16 bits).
5. Virtual gateway is down. Numeric identifier for the virtual gateway (8 bits) and associated adjacent domain (16 bits).
6. Virtual gateway is not reachable according to the given transit policy. Numeric identifier for the virtual gateway (8 bits), associated adjacent domain (16 bits), and transit policy (16 bits).
7. Domain component is not reachable. Numeric identifier for the domain (16 bits) and the component (16 bits).
8. Insufficient resources to establish the path.
9. Target is not reachable via intra-domain routing.
10. No existing path with the given path identifier, in response to a REPAIR message only.

7.6.4. TEARDOWN

The TEARDOWN message type is equal to 3.



PATH ID

(64 bits) Path identifier contained in the original SETUP message.


RSN TYP (8 bits) Numeric identifier for the reason for path teardown.

REASON (variable) Description of the reason for path teardown. Valid reasons include the following types:

1. Virtual gateway is down. Numeric identifier for the virtual gateway (8 bits) and associated adjacent domain (16 bits).
2. Virtual gateway is not reachable according to the given transit policy. Numeric identifier for the virtual gateway (8 bits), associated adjacent domain (16 bits), and transit policy (16 bits).
3. Domain component is not reachable. Numeric identifier for the domain (16 bits) and the component (16 bits).
4. Maximum path lifetime exceeded.
5. Preempted path.
6. Unable to repair path.

7.6.5. ERROR

The ERROR message type is equal to 4.



PATH ID

(64 bits) Path identifier contained in the path control or data message in error.

MSG (8 bits) Numeric identifier for the type of path control message in error. This field is ignored for error type 5.

RSN TYP (8 bits) Numeric identifier for the reason for the PCP message error.

REASON (variable) Description of the reason for the PCP message error. Valid reasons include the following types:

1. Path identifier is already currently active.
2. Domain does not appear in the SETUP message.
3. Transit policy is not configured for the domain. Numeric identifier for the transit policy (16 bits).
4. Virtual gateway not configured for the domain. Numeric identifier for the virtual gateway (8 bits) and associated adjacent domain (16 bits).
5. Unrecognized path identifier in an IDPR data message.

7.6.6. REPAIR

The REPAIR message type is equal to 5. A REPAIR message contains the original SETUP message only.

7.6.7. Negative Acknowledgements

When a policy gateway receives an unacceptable PCP message that passes the CMTTP validation checks, it includes, in its CMTTP ACK, an appropriate negative acknowledgement. This information is placed in the INFORM field of the CMTTP ACK (described previously in section 2.4); the numeric identifier for each type of PCP negative acknowledgement is contained in the left-most 8 bits of the INFORM field. Negative acknowledgements associated with PCP include the following types:

1. Unrecognized PCP message type. Numeric identifier for the unrecognized message type (8 bits).
2. Out-of-date PCP message.
3. Unrecognized path identifier (for all PCP messages except SETUP). Numeric identifier for the unrecognized path (64 bits).

8. Security Considerations

Refer to sections 1.6, 1.7, and 2.3 for details on security in IDPR.

9. Author's Address

Martha Steenstrup
BBN Systems and Technologies
10 Moulton Street
Cambridge, MA 02138

Phone: (617) 873-3192
Email: msteenst@bbn.com

References

- [1] Clark, D., "Policy Routing in Internet Protocols", RFC 1102, May 1989.
- [2] Estrin, D., "Requirements for Policy Based Routing in the Research Internet", RFC 1125, November 1989.
- [3] Little, M., "Goals and Functional Requirements for Inter-Autonomous System Routing", RFC 1126, July 1989.
- [4] Breslau, L. and Estrin, D., "Design of Inter-Administrative Domain Routing Protocols", Proceedings of the ACM SIGCOMM '90 Symposium, September 1990.
- [5] Steenstrup, M., "An Architecture for Inter-Domain Policy Routing", RFC 1478, July 1993.
- [6] Austein, R., "DNS Support for IDPR", Work in Progress, March 1993.
- [7] Bowns, H. and Steenstrup, M., "Inter-Domain Policy Routing Configuration and Usage", Work in Progress, July 1991.
- [8] Woodburn, R., "Definitions of Managed Objects for Inter-Domain Policy Routing (Version 1)", Work in Progress, March 1993.
- [9] McQuillan, J., Richer, I., Rosen, E., and Bertsekas, D., "ARPANET Routing Algorithm Improvements: Second Semiannual Technical Report", BBN Report No. 3940, October 1978.
- [10] Moy, J., "The OSPF Specification", RFC 1131, October 1989.
- [11] Oran, D. (editor), "Intermediate System to Intermediate System Routing Exchange Protocol for Use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", ISO/IEC JTC1/SC6/WG2, October 1989.

- [12] Estrin, D., and Tsudik, G., "Secure Control of Transit Internet-work Traffic, TR-89-15, Computer Science Department, University of Southern California.
- [13] Linn, J., "Privacy Enhancement for Internet Electronic Mail: Part I - Message Encipherment and Authentication Procedures", RFC 1113, August 1989.
- [14] Kent, S., and Linn, J., "Privacy Enhancement for Internet Electronic Mail: Part II - Certificate-based Key Management", RFC 1114, August 1989.
- [15] Linn, J., "Privacy Enhancement for Internet Electronic Mail: Part III - Algorithms, Modes, and Identifiers", RFC 1115, August 1989.
- [16] Rivest, R., "The MD4 Message-Digest Algorithm", RFC 1320, April 1992.
- [17] Rivest, R., "The MD5 Message-Digest Algorithm", RFC 1321, April 1992.