

EXTERIOR GATEWAY PROTOCOL (EGP)

Eric C. Rosen

Bolt Beranek and Newman Inc.

October 1982

It is proposed to establish a standard for Gateway to Gateway procedures that allow the Gateways to be mutually suspicious. This document is a DRAFT for that standard. Your comments are strongly encouraged.

Table of Contents

1 INTRODUCTION..... 1
2 NEIGHBOR ACQUISITION..... 8
3 NEIGHBOR REACHABILITY PROTOCOL..... 11
4 NETWORK REACHABILITY (NR) MESSAGE..... 15
5 POLLING FOR NR MESSAGES..... 22
6 SENDING NR MESSAGES..... 25
7 INDIRECT NEIGHBORS..... 27
8 HOW TO BE A STUB GATEWAY..... 28
9 LIMITATIONS..... 32

1 INTRODUCTION

The DARPA Catenet is expected to be a continuously expanding system, with more and more hosts on more and more networks participating in it. Of course, this will require more and more gateways. In the past, such expansion has taken place in a relatively unstructured manner. New gateways, often containing radically different software than the existing gateways, would be added and would immediately begin participating in the common routing algorithm via the GGP protocol. However, as the internet grows larger and larger, this simple method of expansion becomes less and less feasible. There are a number of reasons for this:

- the overhead of the routing algorithm becomes excessively large;
- the proliferation of radically different gateways participating in a single common routing algorithm makes maintenance and fault isolation nearly impossible, since it becomes impossible to regard the internet as an integrated communications system;
- the gateway software and algorithms, especially the routing algorithm, become too rigid and inflexible, since any proposed change must be made in too many different places and by too many different people.

In the future, the internet is expected to evolve into a set of separate domains or "autonomous systems", each of which consists of a set of one or more relatively homogeneous gateways. The protocols, and in particular the routing algorithm which these gateways use among themselves, will be a private matter, and need never be implemented in gateways outside the particular domain or system.

In the simplest case, an autonomous system might consist of just a single gateway connecting, for example, a local network to the ARPANET. Such a gateway might be called a "stub gateway", since its only purpose is to interface the local network to the rest of the internet, and it is not intended to be used for handling any traffic which neither originated in nor is destined for that particular local network. In the near-term future, we will begin to think of the internet as a set of autonomous systems, one of which consists of the DARPA gateways on ARPANET and SATNET, and the others of which are stub gateways to local networks. The former system, which we shall call the "core" system, will be used as a transport or "long-haul" system by the latter systems.

Ultimately, however, the internet may consist of a number of co-equal autonomous systems, any of which may be used (with certain restrictions which will be discussed later) as a

transport medium for traffic originating in any system and destined for any system. When this more complex configuration comes into being, it will be inappropriate to regard any one autonomous system as a "core" system. For the sake of concreteness, however, and because the initial implementations of the Exterior Gateway Protocol are expected to focus on the case of connecting "stub gateways" to the DARPA gateways on ARPANET and SATNET, we will often use the term "core" gateways in our examples and discussion.

The purpose of the Exterior Gateway Protocol (EGP) is to enable one or more autonomous systems to be used as transport media for traffic originating in some other autonomous system and destined for yet another, while allowing the end-user to see the composite of all the autonomous systems as a single internet, with a flat, uniform address space. The route which a datagram takes through the internet, and the number of autonomous systems which it traverses, are to be transparent to the end-user (unless, of course, the end-user makes use of the IP "source route" option).

In describing the Exterior Gateway Protocol, we have deliberately left a great deal of latitude to the designers and implementers of particular autonomous systems, particularly with regard to timer values. We have done this because we expect that

different gateway implementations and different internet environments may just have different requirements and goals, so that no single strict implementation specification could apply to all. However, this does NOT mean that ANY implementation which conforms to the specification will work well, or that the areas in which we have left latitude are not crucial to performance. The fact that some time-out value, for example, is not specified here does not mean that everything will work no matter what value is assigned.

Autonomous systems will be assigned 16-bit identification numbers (in much the same ways as network and protocol numbers are now assigned), and every EGP message header contains one word for this number. Zero will not be assigned to any autonomous system; rather, the presence of a zero in this field will indicate that no number is present.

We need to introduce the concept of one gateway being a NEIGHBOR of another. In the simplest and most common case, we call two gateways "neighbors" if there is a network to which each has an interface. However, we will need a somewhat more general notion of "neighbor" to allow the following two cases:

- a) Two gateways may be regarded as neighbors if they are directly connected not by a network (in the usual sense

of the term), but by a simple wire, or HDLC line, or some similar means of "direct connection".

- b) Two gateways may be regarded as neighbors if they are connected by an "internet" which is transparent to them. That is, we would like to be able to say that two gateways are neighbors even if they are connected by an internet, as long as the gateways utilize no knowledge of the internal structure of that internet in their own packet-forwarding algorithms.

In order to handle all these cases, let us say that two gateways are NEIGHBORS if they are connected by some communications medium whose internal structure is transparent to them. (See IEN 184 for a more general discussion of this notion of neighbor.)

If two neighbors are part of the same autonomous system, we call them INTERIOR NEIGHBORS; if two neighbors are not part of the same autonomous system, we call them EXTERIOR NEIGHBORS. In order for one system to use another as a transport medium, gateways which are exterior neighbors of each other must be able to find out which networks can be reached through the other. The Exterior Gateway Protocol enables this information to be passed between exterior neighbors. Since it is a polling protocol, it also enables each gateway to control the rate at which it sends

and receives network reachability information, allowing each system to control its own overhead. It also enables each system to have an independent routing algorithm whose operation cannot be disrupted by failures of other systems.

It must be clearly understood that any autonomous system in which routing needs to be performed among gateways within that system must implement its own routing algorithm. (A routing algorithm is not generally necessary for a simple autonomous system which consists of a single stub gateway.) The Exterior Gateway Protocol is NOT a routing algorithm. It enables exterior neighbors to exchange information which is likely to be needed by any routing algorithm, but it does NOT specify what the gateways are to do with this information. The "routing updates" of some autonomous system's interior routing algorithm may or may not be similar in format to the messages of the exterior gateway protocol. The gateways in the DARPA "core" system will initially use the GGP protocol (the old Gateway-Gateway protocol) as their routing algorithm, but this will be subject to change. Gateways in other autonomous systems may use their own Interior Gateway Protocols (IGPs), which may or may not be similar to the IGP of any other autonomous system. They may, of course, use GGP, but will not be permitted to exchange GGP messages with gateways in other autonomous systems.

It must also be clearly understood that the Exterior Gateway Protocol is NOT intended to provide information which could be used as input to a completely general area or hierarchical routing algorithm. It is intended for a set of autonomous systems which are connected in a tree, with no cycles. It does not enable the passing of sufficient information to prevent routing loops if cycles in the topology do exist.

The Exterior Gateway Protocol has three parts: (a) Neighbor Acquisition Protocol, (b) Neighbor Reachability Protocol, and (c) Network Reachability determination. Note that all messages defined by EGP are intended to travel only a single "hop". That is, they originate at one gateway and are sent to a neighboring gateway without the mediation of any intervening gateway. Therefore, the time-to-live field should be set to a very small value. Gateways which encounter EGP messages in their message streams which are not addressed to them may discard them.

2 NEIGHBOR ACQUISITION

Before it is possible to obtain routing information from an exterior gateway, it is necessary to acquire that gateway as a direct neighbor. (The distinction between direct and indirect neighbors will be made in a later section.) In order for two gateways to become direct neighbors, they must be neighbors, in the sense defined above, and they must execute the NEIGHBOR ACQUISITION PROTOCOL, which is simply a standard three-way handshake.

A gateway that wishes to initiate neighbor acquisition with another sends it a Neighbor Acquisition Request. This message should be repeatedly transmitted (at a reasonable rate, perhaps once every 30 seconds or so) until a Neighbor Acquisition Reply is received. The Request will contain an identification number which is copied into the reply so that request and reply can be matched up.

A gateway receiving a Neighbor Acquisition Request must determine whether it wishes to become a direct neighbor of the source of the Request. If not, it may, at its option, respond with a Neighbor Acquisition Refusal message, optionally specifying the reason for refusal. Otherwise, it should send a Neighbor Acquisition Reply message. It must also send a Neighbor

Acquisition Request message, unless it has done so already.

Two gateways become direct neighbors when each has sent a Neighbor Acquisition Message to, and received the corresponding Neighbor Acquisition Reply from, the other.

Unmatched Replies or Refusals should be discarded after a reasonable period of time. However, information about any such unmatched messages may be useful for diagnostic purposes.

A Neighbor Acquisition Message from a gateway which is already a direct neighbor should be responded to with a Reply and a Neighbor Acquisition Message.

If a Neighbor Acquisition Reply is received from a prospective neighbor, but a period of time passes during which no Neighbor Acquisition Message is received from that prospective neighbor, the neighbor acquisition protocol shall be deemed incomplete. A Neighbor Cease message (see below) should then be sent. If one gateway still desires to acquire the other as a neighbor, the protocol must be repeated from the beginning.

If a gateway wishes to cease being a neighbor of a particular exterior gateway, it sends a Neighbor Cease message. A gateway receiving a Neighbor Cease message should always respond with a Neighbor Cease Acknowledgment. It should cease to

treat the sender of the message as a neighbor in any way. Since there is a significant amount of protocol run between direct neighbors (see below), if some gateway no longer needs to be a direct neighbor of some other, it is "polite" to indicate this fact with a Neighbor Cease Message. The Neighbor Cease Message should be retransmitted (up to some number of times) until an acknowledgment for it is received.

Once a Neighbor Cease message has been received, the Neighbor Reachability Protocol (below) should cease to be executed.

NOTE THAT WE HAVE NOT SPECIFIED THE WAY IN WHICH ONE GATEWAY INITIALLY DECIDES THAT IT WANTS TO BECOME A NEIGHBOR OF ANOTHER. While this is hardly a trivial problem, it is not part of the External Gateway Protocol.

3 NEIGHBOR REACHABILITY PROTOCOL

It is important for a gateway to keep real-time information as to the reachability of its neighbors. If a gateway concludes that a particular neighbor cannot be reached, it should cease forwarding traffic to that gateway. To make that determination, a NEIGHBOR REACHABILITY protocol is needed. The EGP protocol provides two messages types for this purpose -- a "Hello" message and an "I Heard You" message.

When a "Hello" message is received from a direct neighbor, an "I Heard You" must be returned to that neighbor "immediately". The delay between receiving a "Hello" and returning an "I Heard You" should never be more than a few seconds.

At the current time, the reachability determination algorithm is left to the designers of a particular gateway. We have in mind algorithms like the following:

A reachable neighbor shall be declared unreachable if, during the time in which we sent our last n "Hello"s, we received fewer than k "I Heard You"s in return. An unreachable neighbor shall be declared reachable if, during the time in which we sent our last m "Hello"s, we received at least j "I Heard You"s in return.

However, the frequency with which the "Hello"s are sent, and the values of the parameters k, n, j, and m cannot be specified here. For best results, this will depend on the characteristics of the neighbor and of the network which the neighbors have in common. THIS IMPLIES THAT THE PROPER PARAMETERS MAY NEED TO BE DETERMINED JOINTLY BY THE DESIGNERS AND IMPLEMENTERS OF THE TWO NEIGHBORING GATEWAYS; choosing algorithms and parameters in isolation, without considering the characteristics of the neighbor and the connecting network, would not be expected to result in optimum reachability determinations.

The "Hello" and "I Heard You" messages have a status field which the sending gateway uses to indicate whether it thinks the receiving gateway is reachable or not. This information can be useful for diagnostic purposes. It also allows one gateway to make its reachability determination parasitic on the other: only one gateway actually needs to send "Hello" messages, and the other can declare it up or down based on the status field in the "Hello". That is, the "passive" gateway (which sends only "I Heard You"s) declares the "active" one (which sends only "Hello"s) to be reachable when the "Hello"s from the active one indicate that it has declared the passive one to be reachable. Of course, this can only work if there is prior agreement as to which neighbor is to be the active one. (Ways of coming to this

"prior agreement" are not part of the Exterior Gateway Protocol.)

A direct neighbor gateway should also be declared unreachable if the network connecting it supplies lower level protocol information from which this can be deduced. Thus, for example, if a gateway receives an 1822 Destination Dead message from the ARPANET which indicates that a direct neighbor is dead, it should declare that neighbor unreachable. The neighbor should not be declared reachable again until the requisite number of Hello/I-Heard-You packets have been exchanged.

A direct neighbor which has become unreachable does not thereby cease to be a direct neighbor. The neighbor can be declared reachable again without any need to go through the neighbor acquisition protocol again. However, if the neighbor remains unreachable for an extremely long period of time, such as an hour, the gateway should cease to treat it as a neighbor, i.e., should cease sending Hello messages to it. The neighbor acquisition protocol would then need to be repeated before it could become a direct neighbor again.

"Hello" and "I Heard You" messages from gateway G to gateway G' also carry the identification number of the NR poll message (see below) which G has most recently received from G'.

"Hello" and "I Heard You" messages from gateway G to gateway G' also carry the minimum interval in minutes with which G is willing to be polled by G' for NR messages (see below).

"Hello" messages from sources other than direct neighbors should simply be ignored. However, logging the presence of any such messages might provide useful diagnostic information.

A gateway which is going down, or whose interface to the network which connects it to a particular neighbor is going down, should send a Gateway Going Down message to all direct neighbors which will no longer be able to reach it. It should retransmit that message (up to some number of times) until it receives a Gateway Going Down Acknowledgment. This provides the neighbors with an advance warning of an outage, and enables them to prepare for it in a way which will minimize disruption to existing traffic.

4 NETWORK REACHABILITY (NR) MESSAGE

Terminology: Let gateway G have an interface to network N. We say that G is AN APPROPRIATE FIRST HOP to network M relative to network N (where M and N are distinct networks) if and only if the following condition holds:

Traffic which is destined for network M, and which arrives at gateway G over its network N interface, will be forwarded to M by G over a path which does not include any other gateway with an interface to network N.

In short, G is an appropriate first hop for network M relative to network N just in case there is no better gateway on network N through which to route traffic which is destined for network M. For optimal routing, traffic in network N which is destined for network M ought always to be forwarded to a gateway which is an appropriate first hop.

In order for exterior neighbors G and G' (which are neighbors over network N) to be able to use each other as packet switches for forwarding traffic to remote networks, each needs to know the list of networks for which the other is an appropriate first hop. The Exterior Gateway Protocol defines a message, called the Network Reachability Message (or NR message), for transferring this information.

Let G be a gateway on network N. Then the NR message which G sends about network N must contain the following information:

A list of all the networks for which G is an appropriate first hop relative to network N.

If G' can obtain this information from exterior neighbor G, then it knows that no traffic destined for networks which are NOT in that list should be forwarded to G. (It cannot simply conclude, however, that all traffic for any networks in that list ought to be forwarded via G, since G' may also have other neighbors which are also appropriate first hops to network N. For example, G and G'' might each be neighbors of G', but might be "equidistant" from some network M. Then each could be an appropriate first hop.)

For each network in the list, the NR message also contains a byte which specifies the "distance" (according to some metric whose definition is left to the designers of the autonomous system of which gateway G is a member) from G to that network. This information might (or might not) be useful in the interior routing algorithm of gateway G', or for diagnostic purposes.

The maximum value of distance (255.) shall be taken to mean that the network is UNREACHABLE. ALL OTHER VALUES WILL BE TAKEN TO MEAN THAT THE NETWORK IS REACHABLE.

If an NR message from some gateway G fails to mention some network N which was mentioned in the previous NR message from G, it shall be assumed that N is still reachable from G. HOWEVER, IF N IS NOT MENTIONED IN TWO SUCCESSIVE NR MESSAGES FROM G, THAT SHALL BE TAKEN TO MEAN THAT N IS NO LONGER REACHABLE FROM G. This procedure is necessary to ensure that networks which can no longer be reached, but which are never explicitly declared unreachable, are timed out and removed from the list of reachable networks.

It may often be the case that where G and G' are exterior neighbors on network N, G knows of many more gateway neighbors on network N, and knows for which networks those other neighbors are the appropriate first hop. Since G' may not know about all these other neighbors, it is convenient and often more efficient for it to be able to obtain this information from G. Therefore, the EGP NR message also contains fields which allow G to specify the following information:

- a) A list of all neighbors (both interior and exterior) of G (on network N) which G has reliably determined to be reachable. Gateways should be included in this list only if G is actively running its neighbor reachability protocol with them.

b) For each of those neighbors, the list of networks for which that neighbor is an appropriate first hop (relative to network N).

c) For each such <neighbor, network> pair, the "distance" from that neighbor to that network.

Thus the NR message provides a means of allowing a gateway to "discover" new neighbors by seeing whether a neighbor that it already knows of has any additional neighbors on the same network. This information also makes possible the implementation of the INDIRECT NEIGHBOR strategy defined below.

A more precise description of the NR message is the following.

The data portion of the message will consist largely of blocks of data. Each block will be headed by a gateway address, which will be the address either of the gateway sending the message or of one of that gateway's neighbors. Each gateway address will be followed by a list of the networks for which that gateway is an appropriate first hop, and the distance from that gateway to each network.

Preceding the list of data blocks is:

a) The address of the network which this message is about.

If G and G' are neighbors on network N, then in the NR message going from G to G', this is the address of network N. For convenience, four bytes have been allocated for this address -- the trailing one, two, or three bytes should be zero.

- b) The count (one byte) of the number of interior neighbors of G for which this message contains data blocks. By convention, this count will include the data block for G itself, which should be the first one to appear.
- c) The count (one byte) of the number of exterior neighbors of G for which this message contains data blocks.

Then follow the data blocks themselves, first the block for G itself, then the blocks for all the interior neighbors of G (if any), then the blocks for the exterior neighbors. Since all gateways mentioned are on the same network, whose address has already been given, the gateway addresses are given with the network address part (one, two, or three bytes) omitted, to save space.

Each block includes a one-byte count of the number of networks for which that gateway is the appropriate first hop. In the list of networks, each network address is either one, two, or three bytes, depending on whether it is a class A, class B, or

class C network. No trailing bytes are used.

It may sometimes be necessary to fragment the NR message. The NR message contains a byte indicating the number of this fragment (fragments will be numbered from zero), and a byte containing the number of the last fragment (NOT the number of fragments). If fragmentation is not used, these bytes must both be zero. EACH FRAGMENT MUST BE A FULLY SELF-CONTAINED NR MESSAGE. That is, each fragment will begin with a count of interior and exterior neighbors, and will have some integral number of gateway data blocks. The number of data blocks in each fragment must correspond to the neighbor counts at the beginning of that fragment. However, only the first fragment should begin with a data block describing the sending gateway.

This scheme enables each fragment to be processed independently, and requires no complex reassembly mechanisms. It also enables processing of a message all of whose fragments have not been received. If, after some amount of time and some number of retransmissions of a poll, not all fragments have been received, the fragments which are present shall be processed as if they constituted the complete NR message. (This means that networks mentioned only in the missing fragment will retain the "distance" values they had in the previous NR message from that gateway. However, if no new value for a particular network is

received in the next NR message from that gateway, the network will be declared unreachable.)

5 POLLING FOR NR MESSAGES

No gateway is required to send NR messages to any other gateway, except as a response to an NR Poll from a direct neighbor. However, a gateway is required to respond to an NR Poll from a direct neighbor within several seconds (subject to the qualification two paragraphs hence), even if the gateway believes that neighbor to be down.

The EGP NR Poll message is defined for this purpose. No gateway may poll another for an NR message more often than once per minute. A gateway receiving more than one poll per minute may simply ignore the excess polls, or may return an error message. The Hello and I Heard You messages which gateway G sends to gateway G' indicate the minimum interval which G will accept as the polling interval from G'. That is, G' will not guarantee to respond to polls from G that arrive less than that interval apart.

Polls must only be sent to direct neighbors which are declared reachable by the neighbor reachability protocol.

An NR Poll message contains an identification number chosen by the polling gateway. The polled gateway will return this number in the NR message it sends in response to the poll, to enable the polling gateway to match up received NR messages with

polls. It will be the responsibility of the polling gateway to choose an identification number which is sufficiently "unique" to allow detection of out-of-date NR messages which may still be floating around the network. Since polls are relatively infrequent, this is not expected to be much of a problem. However, to aid in choosing an identification number, the Hello and I Heard You messages carry the identification number of the last NR poll received from the neighbor to which they are being sent.

In general, a poll should be retransmitted some number of times (with a reasonable interval between retransmissions) until an NR message is received. IF NO NR MESSAGE IS RECEIVED AFTER THE MAXIMUM NUMBER OF RETRANSMISSIONS, THE POLLING GATEWAY SHOULD ASSUME THAT THE POLLED GATEWAY IS NOT AN APPROPRIATE FIRST HOP FOR ANY NETWORK WHATSOEVER. The optimum parameters for the polling/retransmission algorithm will be dependent on the characteristics of the two neighbors and of the network connecting them.

If only some fragments of an NR message are received after the maximum number of retransmissions, the fragments that are present shall be treated as constituting the whole of the NR message.

Received NR messages whose identification numbers do not match the identification number of the most recently sent poll shall be ignored. There is no provision for multiple outstanding polls to the same neighbor.

6 SENDING NR MESSAGES

In general, NR messages are to be sent only in response to a poll. However, between two successive polls from an exterior neighbor, a gateway may send one and only one unsolicited NR message to that neighbor. This gives it limited ability to quickly announce network reachability changes that may have occurred in the interval since the last poll. Excess unsolicited NR messages may be ignored, or an error message may be returned.

An NR message should be sent within several seconds after receipt of a poll. Failure to respond in a timely manner to an NR poll may result in the polling gateway's deciding that the polled gateway is not an appropriate first hop to any network.

NR messages sent in response to polls carry the identification number of the poll message in their "identification number" fields. Unsolicited NR messages carry the identification number of the last poll received, and have the "unsolicited" bit set. (Note that this allows for only a single unsolicited NR message per polling period.)

To facilitate the sending of unsolicited NR messages, the NR poll message has a byte indicating the polling interval in minutes.

Polls from non-neighbors, from neighbors which are not declared reachable, or with bad IP source network fields, should be responded to with an EGP error message with the appropriate "reason" field. If G sends an NR poll to G' with IP source network N, and G' is not a neighbor of G on its interface to network N (or G' does not have an interface to network N), then the source network field is considered "bad".

Duplicated polls (successive polls with the same identification number) should be responded to with duplicates of the same NR message. If that message is fragmented, the same fragments shall be sent each time. Note that there is no provision for handling multiple outstanding polls from a single neighbor. NOTE THAT IF THE SAME FRAGMENTS ARE NOT SENT IN RESPONSE TO DUPLICATED POLLS, INCORRECT REASSEMBLY WILL BE THE PROBABLE RESULT. If fragmentation is not being used, however, then no harm should result from responding to a duplicate poll with a different (presumably more recent) NR message.

7 INDIRECT NEIGHBORS

Becoming a "direct neighbor" of an exterior gateway requires three steps: (a) neighbor acquisition, (b) running a neighbor reachability protocol, and (c) polling the neighbor periodically for NR messages. Suppose, however, that gateway G receives an NR message from G', in which G' indicates the presence of other neighbors G1, ..., Gn, each of which is an appropriate first hop for some set of networks to which G' itself is not an appropriate first hop. Then G should be allowed to forward traffic for those networks directly to the appropriate one of G1, ..., Gn, without having to send it to G' first. In this case, G may be considered an INDIRECT NEIGHBOR of G1, ..., Gn, since it is a neighbor of these other gateways for the purpose of forwarding traffic, but does not perform neighbor acquisition, neighbor reachability, or exchange of NR messages with them. Neighbor and network reachability information is obtained indirectly via G', hence the designation "indirect neighbor". We say that G is an indirect neighbor of G1, ..., Gn VIA G'.

If G is an indirect neighbor of G' via G'', and then G receives an NR message from G'' which does not mention G', G should treat G' as having become unreachable.

8 HOW TO BE A STUB GATEWAY

The most common application of EGP will probably be its use to enable a stub gateway to communicate with one of the DARPA core gateways, so as to enable data flow between networks accessible only via the stub and networks accessible only via the system of core gateways. As discussed previously, a stub gateway can be considered to be a one-gateway internet system with no interior neighbors. It is probably used to interface a local network or networks to a long range transport network (such as ARPANET or SATNET) on which there is a core gateway. In this case, the stub will not want the core gateways to forward it any traffic other than traffic which is destined for the network or networks which can be reached only via the stub. In general, the stub will not want to perform any services for the internet transport system which are not needed in order to be able to pass traffic to and from the networks that cannot be otherwise reached.

The stub should have tables configured in with the addresses of a small number of the core gateways (no more than two or three) with which it has a common network. It will be the responsibility of the stub to initiate neighbor acquisition with these gateways. When a stub and a core gateway become direct neighbors, the core gateway will begin sending Hello messages.

When the stub declares the core gateways which are direct neighbors to be reachable, it should poll those gateways for NR messages at a rate not to exceed once per minute (or as specified in the Hello messages from the core gateways). The core gateways will also poll the stub for NR messages.

The NR message sent by the stub should be the simplest allowable. That is, it should have only a single data block, headed by its own address (on the network it has in common with the neighboring core gateway), listing just the networks to which it is an appropriate first hop. These will be just the networks that can be reached no other way, in general.

The core gateways will send complete NR messages, containing information about all other gateways on the common networks, both core gateways (which shall be listed as interior neighbors) and other gateways (which shall be listed as exterior neighbors, and may include the stub itself). This information will enable the stub to become an indirect neighbor of all these other gateways. That is, the stub shall forward traffic directly to these other gateways as appropriate, but shall not become direct neighbors with them.

The core gateways will report distances less than 128 if the network can be reached without leaving the core system (i.e.,

without traversing any gateway other than a core gateway), and greater than or equal to 128 otherwise.

The stub should NEVER forward to any (directly or indirectly) neighboring core gateway any traffic for which that gateway is not an appropriate first hop, as indicated in an NR message. Of course, this does not apply to datagrams which are using the source route option; any such datagrams should always be forwarded as indicated in the source route option field, even if that requires forwarding to a gateway which is not an appropriate first hop.

If the direct neighbors of a stub should all fail, it will be the responsibility of the stub to acquire at least one new direct neighbor. It can do so by choosing one of the core gateways which it has had as an indirect neighbor, and executing the neighbor acquisition protocol with it. (It is possible that no more than one core gateway will ever agree to become a direct neighbor with any given stub gateway at any one time.)

If the stub gateway does not respond in a timely manner to Hello messages from the core gateway, it may be declared unreachable. If it does not respond to NR poll messages in a timely manner, its networks may be declared unreachable. In both these cases, the core gateways may discard traffic destined for

those networks, returning ICMP "destination network unreachable" to the source hosts.


The stub gateway is expected to fully execute the ICMP protocol, as well as the EGP protocol. In particular, it must respond to ICMP echo requests, and must send ICMP destination dead messages as appropriate. It is also required to send ICMP Redirect messages as appropriate.

9 LIMITATIONS

It must be clearly understood that the Exterior Gateway Protocol does not in itself constitute a network routing algorithm. In addition, it does not provide all the information needed to implement a general area routing algorithm. If the topology of the set of autonomous systems is not tree-structured (i.e., if it has cycles), the Exterior Gateway Protocol does not provide enough topological information to prevent loops.

If any gateway sends an NR message with false information, claiming to be an appropriate first hop to a network which it in fact cannot even reach, traffic destined to that network may never be delivered. Implementers must bear this in mind.

NEIGHBOR ACQUISITION MESSAGE



Description:

The Neighbor Acquisition messages are used by interior and exterior gateways to become neighbors of each other.

EGP Version #

1

Type

3

Code

- Code = 0 Neighbor Acquisition Request
- Code = 1 Neighbor Acquisition Reply
- Code = 2 Neighbor Acquisition Refusal (see Info field)
- Code = 3 Neighbor Cease Message (see Info field)
- Code = 4 Neighbor Cease Acknowledgment

Checksum

The EGP checksum is the 16-bit one's complement of the one's complement sum of the EGP message starting with the EGP version number field. For computing the checksum, the checksum field should be zero.

Autonomous System #

This 16-bit number identifies the autonomous system containing the gateway which is the source of this message.

Info

For Refusal message, gives reason for refusal:

- 0 Unspecified
- 1 Out of table space
- 2 Administrative prohibition

For Cease message, gives reason for ceasing to be neighbor:

- 0 Unspecified
- 1 Going down
- 2 No longer needed

Otherwise, this field MUST be zero.

Identification Number

An identification number to aid in matching requests and replies.

NEIGHBOR HELLO/I HEARD YOU MESSAGE

```

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
+-----+-----+-----+-----+-----+-----+-----+-----+
! EGP Version # !   Type           !   Code           !   Status         !
+-----+-----+-----+-----+-----+-----+-----+-----+
!   Checksum           !   Autonomous System #           !
+-----+-----+-----+-----+-----+-----+-----+-----+
!   Sequence #           !Min Poll Intvl !   Zero           !
+-----+-----+-----+-----+-----+-----+-----+-----+
!   Last Poll Id #           !
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Description:

Exterior neighbors use EGP Neighbor Hello and I Heard You Messages to determine neighbor connectivity. When a gateway receives an EGP Neighbor Hello message from a neighbor it should respond with an EGP I Heard You message.

EGP Version #

1

Type

5

Code

Code = 0 for Hello
Code = 1 for I Heard you

Checksum

The EGP checksum is the 16-bit one's complement of the one's complement sum of the EGP message starting with the EGP version number field. For computing the checksum, the checksum field should be zero.

Autonomous System #

This 16-bit number identifies the autonomous system containing the gateway which is the source of this message.

Sequence Number

A sequence number to aid in matching requests and replies.

Status

- 0 No status given
- 1 You appear reachable to me
- 2 You appear unreachable to me due to neighbor reachability protocol
- 3 You appear unreachable to me due to network reachability information (such as 1822 "destination dead" messages from ARPANET)
- 4 You appear unreachable to me due to problems with my network interface

Last Poll Id Number

The identification number of the most recently received NR poll message from the neighbor to which this message is being sent.

Minimum Polling Interval

This gateway should not be polled for NR messages more often than once in this number of minutes.

NR POLL Message

```

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
+-----+-----+-----+-----+-----+-----+-----+-----+
! EGP Version # !      Type      !      Code      !      Unused      !
+-----+-----+-----+-----+-----+-----+-----+-----+
!   Checksum           !           Autonomous System #           !
+-----+-----+-----+-----+-----+-----+-----+-----+
!           IP Source Network           ! Interval           !
+-----+-----+-----+-----+-----+-----+-----+-----+
!   Identification #           !
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Description:

A gateway that wants to receive an NR message from an Exterior Gateway will send an NR Poll message. Each gateway mentioned in the NR message will have an interface on the network that is in the IP source network field.

EGP Version #

1

Type

2

Code

0

Checksum

The EGP checksum is the 16-bit one's complement of the one's complement sum of the EGP message starting with the EGP version number field. For computing the checksum, the checksum field should be zero.

Autonomous System #

This 16-bit number identifies the autonomous system containing the gateway which is the source of this message.

Identification Number

An identification number to aid in matching requests and replies.

IP Source Network

Each gateway mentioned in the NR message will have an interface on the network that is in the IP source network field. The IP source network is coded as one byte of network number followed by two bytes of zero for class A networks, two bytes of network number followed by one byte of zero for class B networks, and three bytes of network number for class C networks.

Interval

The polling interval in minutes.

NETWORK REACHABILITY MESSAGE

```

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
+-----+-----+-----+-----+
! EGP Version # !      Type      !   Code      !U! Zeroes  !
+-----+-----+-----+-----+
!   Checksum    !           !   Autonomous System #   !
+-----+-----+-----+-----+
! Fragment #   !# of last frg. !   Identification #   !
+-----+-----+-----+-----+
!                               IP Source Network                               !
+-----+-----+-----+-----+
! # of Int Gwys ! # of Ext Gwys !
+-----+-----+-----+-----+
! # of Nets    !                               ; # of nets for
+-----+-----+-----+-----+ Gateway 1
! Gateway 1 IP address (without network #) ! ; 1, 2 or 3 bytes
+-----+-----+-----+-----+
! net 1,1      !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!! ; 1, 2 or 3 bytes
+-----+-----+-----+-----+
! distance    !
+-----+-----+-----+-----+
! net 1,2      !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!! ; 1, 2 or 3 bytes
+-----+-----+-----+-----+
! distance    !
+-----+-----+-----+-----+
.
.
+-----+-----+-----+-----+
! net 1,m      !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!! ; m nets reachable
+-----+-----+-----+-----+ ; via Gateway 1
.
.
+-----+-----+-----+-----+
! # of nets    !           ;number of nets for Gateway n
+-----+-----+-----+-----+
!           Gateway n IP address (without network #)           !
+-----+-----+-----+-----+
! net n,1      !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!! ; 1, 2 or 3 bytes
+-----+-----+-----+-----+
! distance    !
+-----+-----+-----+-----+

```

```

+-----+
! net n,2      !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!! ; 1, 2 or 3 bytes
+-----+
! distance    !      .
+-----+

```

```

+-----+
! net n,m      !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!! ; m nets reachable
+-----+ ; via Gateway n
! distance    !
+-----+

```

Description:

The Network Reachability message (NR) is used to discover which networks may be reached through Exterior Gateways. The NR message is sent in response to an NR Poll message.

EGP Version

1

Type

1

Code

0

Checksum

The EGP checksum is the 16-bit one's complement of the one's complement sum of the EGP message starting with the EGP version number field. For computing the checksum, the checksum field should be zero.

Autonomous System

This 16-bit number identifies the autonomous system containing the gateway which is the source of this message.

U (Unsolicited) bit

This bit is set if the NR message is being sent unsolicited.

Identification Number

The identification number of the last NR poll message received from the neighbor to whom this NR message is being sent. This number is used to aid in matching polls and replies.

Fragment Number

Which Fragment this is in the NR Message. Zero, if fragmentation is not used.

Number of Last Fragment

Number of the last fragment in the NR Message. Zero, if fragmentation is not used.

IP Source Network

Each gateway mentioned in the NR message will have an interface on the network that is in the IP source network field.

of Interior Gateways

The number of interior gateways that are mentioned in this message.

of Exterior Gateways

The number of exterior gateways that are mentioned in this message.

of Networks

The number of networks for which the gateway whose IP address immediately follows is the appropriate first hop.

Gateway IP address

1, 2 or 3 bytes of Gateway IP address (without network #).

Network address

1, 2, or 3 bytes of network address of network which can be reached via the preceding gateway.

Distance

1 byte of distance in # of hops.

EGP ERROR MESSAGE

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
! EGP Version #										! Type										! Code										! Unused									
! Checksum										! Autonomous System #																													
! Error Type										! Error Code										! Id. # of Erroneous Msg.																			
! Sequence #																																							

Description:

An EGP Error Message is sent in response to an EGP Message that has a bad checksum or has an incorrect value in one of its fields.

EGP Version #

1

Type

8

Code

0

Checksum

The EGP checksum is the 16-bit one's complement of the one's complement sum of the EGP message starting with the EGP version number field. For computing the checksum, the checksum field should be zero.

Autonomous System #

This 16-bit number identifies the autonomous system containing the gateway which is the source of this message.

Sequence Number

A sequence number assigned by the gateway sending the error message.

Error Type

The Type of the EGP message that was in error.

Error Code

The Code of the EGP message that was in error.

Identification number of erroneous message

The Sequence number of the EGP message that was in error.

Reason

The reason that the EGP message was in error. The following reasons are defined:

- 0 - unspecified
- 1 - Bad EGP checksum
- 2 - Bad IP Source address in NR Poll or Response
- 3 - Undefined EGP Type or Code
- 4 - Received poll from non-neighbor
- 5 - Received excess unsolicited NR message
- 6 - Received excess poll
- 7 - Erroneous counts in received NR message
- 8 - No response received to NR poll
- 9 - Not all fragments of NR message received

