

IDRP for IP

Status of this memo

This memo specifies the adaptation of the ISO Inter-Domain Routing Protocol ([1]) that enables it to be used as an Inter-Autonomous System routing protocol in the TCP/IP Internet. IDRP with this adaptation will be called "IDRP for IP" in this document. Dual IDRP, that is, a single instance of IDRP that can simultaneously support Inter-Domain/Inter-Autonomous System routing in the TCP/IP and OSI Internets is outside the scope of this document.

This document is an Internet Draft. Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its Areas, and its Working Groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six months. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "work in progress."

[1](#). Overview

IDRP ([1]) is defined as the protocol for exchange of Inter-Domain routing information between routers to support forwarding of ISO 8473 (Connectionless Network Layer Protocol (CLNP)) [2] PDUs.

The network reachability information exchanged via IDRP provides sufficient information to detect routing loops and enforce routing decisions based on performance preference and policy constraints as outlined in [RFC 1104](#) [9]. In particular, IDRP exchanges routing information containing full domain-level paths and enforces routing policies based on configuration information.

As the Internet has evolved and grown over in recent years, it has become evident that it is soon to face several serious scaling problems. These include:

Exhaustion of the class B network address space. One fundamental

Expiration Date February 1994

[Page 1]

INTERNET DRAFT

September 1993

cause of this problem is the lack of a network class of a size which is appropriate for mid-sized organization; class C, with a maximum of 254 host addresses, is too small while class B, which allows up to 65534 addresses, is too large to be densely populated. Growth of routing tables in Internet routers are beyond the ability of current software (and people) to effectively manage. Eventual exhaustion of the 32-bit IP address space.

It has become clear that the first two of these problems are likely to become critical within the next one to three years. Classless inter-domain routing (CIDR) [7], [10] attempts to deal with these problems by proposing a mechanism to slow the growth of the routing table and the need for allocating new IP network numbers. It does not attempt to solve the third problem, which is of a more long-term nature, but instead endeavors to ease enough of the short to mid-term difficulties to allow the Internet to continue to function efficiently while progress is made on a longer-term solution.

IDRP may be viewed as an extension of BGP-4 [11] that provides (among other things) much better scaling with respect to support for routing information aggregation and reduction based on CIDR, as well as stronger capabilities for policy based routing (e.g. ability to impose control over transit traffic).

This document specifies the appropriate adaptation of the IDRP protocol definition that enables it to be used as a protocol for the exchange of inter-autonomous system information among routers to support the forwarding of IP packets across multiple autonomous systems.

The adaptation is defined in such a way that a Dual IDRP will be able to fully interoperate with IDRP for IP.

[2.](#) Terminology

This document assumes that the reader is familiar with the following documents:

IP protocol specification ([RFC 791](#))[5], and IDRP specification (IS 10747).

A few definitions are in order to aid the reader:

BIS - a Boundary Intermediate System (or border router)

BISPDU - an IDRP message exchanged between a pair of BISs

FIB - Forwarding Information Base (IP forwarding table)

IS - Intermediate System (router)

NET - Network Entity Title - an ISO network address for a router

Expiration Date February 1994

[Page 2]

INTERNET DRAFT

September 1993

NLRI - Network Layer Reachability Information (set of reachable destinations)

NPDU - an IP packet

PDU - a packet

SNPA - subnetwork point of attachment (MAC address)

[3.](#) Assumptions

The IDRP for IP protocol assumes that within a single connected internet network addresses are unique. The IDRP for IP protocol cannot be guaranteed to work in an environment where network addresses within a single connected internet are not unique.

All of the discussions in this document are based on the assumption that the Internet is a collection of arbitrarily connected Autonomous Systems. That is, the Internet will be modeled as a general graph whose nodes are AS's and whose edges are connections between pairs of AS's.

The classic definition of an Autonomous System is a set of routers

under a single technical administration, using an interior gateway protocol and common metrics to route packets within the AS and using an exterior gateway protocol to route packets to other AS's. Since this classic definition was developed, it has become common for a single AS to use several interior gateway protocols and sometimes several sets of metrics within an AS. The use of the term Autonomous System here stresses the fact that, even when multiple IGPs and metrics are used, the administration of an AS appears to other AS's to have a single coherent interior routing plan and presents a consistent picture of which networks are reachable through it.

AS's are assumed to be administered by a single administrative entity, at least for the purposes of representation of routing information to systems outside of the AS.

4. The Adaptation Layer

The Inter-Domain Routing Protocol (IDRP) or, more formally,

"The Protocol for the Exchange of Inter-Domain Routing information among Intermediate Systems to support Forwarding of ISO 8473 PDUs (IDRP)"

is the inter-domain routing protocol defined to support the forwarding of Connectionless Network Layer Protocol (CLNP) [ISO 8473] packets that traverse multiple routing domains.

Expiration Date February 1994

[Page 3]

INTERNET DRAFT

September 1993

While this protocol was developed within ISO, it makes few, if any, ISO-specific assumptions. In particular, it does not require participating domains to support any specific ISO Intra-Domain protocol, such as IS-IS (ISO IS 10589) [3], nor does it require participating routers to run ES-IS (ISO 9542) [4]. The only requirements imposed by the protocol on the participating routers is that the protocol information can be exchanged between them over a connectionless network layer, which in the case of OSI is CLNP, and that the network layer connectivity between routers within a single routing domain should be provided by means outside of IDRP (e.g., via some intra-domain routing protocol). IDRP does not place any restrictions on the structure of reachability information, as long it can be expressed as an arbitrary set of variable length address prefixes.

Since IP can provide connectionless service between routers, and since reachable IP destinations can be expressed as IP address prefixes, IDRP can be easily adapted to be an Inter-Autonomous System routing protocol which can be used in the pure TCP/IP Internet.

While conceptually it is possible to define a mapping between the security field of an IP header and IDRP NPDU-derived distinguishing attributes, this mapping is outside the scope of this document. In addition, address-specific QoSs (Source Specific QoS and Destination Specific QoS) have no counterparts in IP. Therefore, the use of the following IDRP distinguishing attributes for IP packets will not be defined in this document: Priority Locally Defined QoS Security

Mapping between the following IDRP distinguishing attributes: Transit Delay Residual Error Expense

and the IP Type of Service (TOS) parameters is defined in [Section 5.2.3](#) of this document.

Note that an implementation that does not support any of the NPDU-derived distinguishing attributes can fully interoperate with an implementation that does support them. Therefore, an IDRP for IP implementation that will support routing sensitive to the parameters present in the TOS field of the IP header will be compatible with the implementation that does not provide such support.

[5](#). Implementor's Guide for IP specific functions.

In order to implement IDRP for IP, only a subset of the features of the IDRP protocol must be implemented.

[5.1](#) Features in IDRP which need not be implemented

The functions of the IDRP protocol which shall not be implemented for IDRP for IP are those functions which deal with the following (all

Locally Defined QoS according to [section 7.12.11](#) Security according to [section 7.12.14](#) Priority according to [section 7.12.16](#) Forwarding CLNP packets according to [section 8](#) The interface to CLNP according to [section 9](#) support of the Network Management information described in the IDRП GDMO according to [section 11](#)

Therefore, with IDRП for IP the following items dealing with CLNP in the IDRП conformance clause (section 12.1 of [1]) shall not be implemented:

clause (d): Locally Defined QoS, Security, Priority clause (r) clause (s) clause (t)

[5.1.1](#) PICS Table Information

The PICS (Protocol Implementation Conformance Statement) provides a convenient and concise mechanism to define which functions need and need not be implemented for IDRП for IP. All references in this section are with respect to [1]. All items with PICS Status as Optional need not be implemented in IDRП for IP. Specifically, IDRП for IP does not require support for the following items:

A.4.3 Table 9:

MGT

A.4.8 (Table 14):

PSRCRT, DATTS, MATCH

A.4.11 (Table 17):

LQOSG, SECG, PRTY

A.4.11 (Table 18):

LQOSP, SECP, PRTYP

A.4.11 (Table 19):

LQOSR, SECR, PRTYR

Implementation of all other items with Optional Status not listed in the previous paragraph is optional.

[5.2](#) Features in IDRП which shall be implemented

An implementation of IDRП for IP shall contain all mandatory features of IDRП, except those mentioned in [Section 5.1](#) of this document. In addition, a BIS for IDRП for IP shall implement:

an interface to the IP protocol described in [section 5.2.1](#) of this

Expiration Date February 1994

[Page 5]

INTERNET DRAFT

September 1993

document the ability to identify and extract IP reachability and IP forwarding information as described in [section 5.2.2](#) of this document IP packet forwarding functions described in [section 5.2.9](#) of this document domain configuration information listed in [section 5.2.8](#) of this document the advertisement of IP address information in NLRI as described in [section 5.3](#) of this document

[5.2.1](#) Exchanging IDRP information between IP-only routers.

IDRP assumes pair-wise communication between participating BISs. IDRP information is carried between a pair of BISs in the form of BISPDU's. In the case of IDRP for IP, these BISPDU's are carried in the data field of IP packets of protocol type 45.

[5.2.2](#) Identifying IP reachability and IP forwarding information

NLRI passed by the UPDATE PDU has an indication of protocol type. An implementation of IDRP for IP shall have the following values in the NLRI field:

Proto_Type: 1 (ISO/IEC TR 9577 IPI/SPI)


Proto_Length: 1

Protocol: hexadecimal 'CC'

Addr_Length: variable (the value shall be between 4 and 32)

Addr_Info: IP address prefixes, encoded in binary, as follows:

This is a variable length field that contains a list of IP address prefixes for the routes that are being advertised. Each IP address prefix is encoded as a 2-tuple of the form <length, prefix>, whose fields are described below:



The use and the meaning of these fields are as follows:

a) Length:

The Length field indicates the length in bits of the IP

Expiration Date February 1994

[Page 6]

INTERNET DRAFT

September 1993

address prefix. A length of zero indicates a prefix that matches all IP addresses (with prefix, itself, of zero octets).

b) Prefix:

The Prefix field contains IP address prefixes followed by enough trailing bits to make the end of the field fall on an octet boundary. Note that the value of trailing bits is irrelevant.

An implementation of IDRP for IP shall ignore any NLRI indicating a different protocol type.

The NEXT_HOP attribute in the UPDATE PDU also has a Type field which indicates the protocol family for the address of the NEXT_HOP. An implementation of IDRP for IP shall have the following values in the NEXT_HOP field:

Proto_Type: 1 (ISO/IEC TR 9577 IPI/SPI)

Proto_Length: 1

Protocol: hexadecimal 'CC'

Length of NET: 4

NET of Next Hop: an IP address, encoded in binary as specified in

[section 5.2.2.](#)

SNPA information: as appropriate for the subnetwork type in use

An implementation of IDRP for IP shall ignore any NEXT_HOP information indicating a different protocol type.

[5.2.3](#) Mapping between IP Type Of Service parameters and IDRP distinguishing attributes.

IP defines four distinct Type of Service (TOS) Parameters (see [\[12\]](#)):
minimize delay maximize throughput maximize reliability minimize monetary cost

An IP packet can carry at most one of the above TOSs. Therefore, a RIB in IDRP for IP can have at most one distinguishing attribute.

IDRP for IP supports the following distinguishing attributes: Transit Delay Residual Error Expense

A conformant implementation is required to recognize these attributes when received from an adjacent BIS.

An IP-derived distinguishing attribute is defined as the TOS parameter extracted from an IP packet that needs to be forwarded by a BIS.

The mapping between the IP-derived distinguishing attribute and a RIB-Att is defined as follows:

IP TOS	IDRP distinguishing attribute
-----	-----
minimize delay	Transit Delay
maximize reliability	Residual Error
minimize monetary cost	Expense

[5.2.4](#) ATOMIC_AGGREGATE Attribute

A new optional transitive attribute, ATOMIC_AGGREGATE, is defined for use in IDRP for IP. This attribute is intended to facilitate interoperation between IDRP for IP and BGP-4. The type code of this attribute shall be 129. This attribute has a length of zero octets.

This attribute shall be handled as follows:

Any IDRP for IP router receiving a route with the ATOMIC_AGGREGATE attribute shall not deaggregate that route.

[5.2.5](#) AGGREGATE Attribute

A new optional transitive attribute, AGGREGATOR is defined for use in IDRP for IP. This attribute is intended to facilitate interoperation between IDRP for IP and BGP-4. The type code of this attribute shall be 130. This attribute shall have the following encoding:

proto_type - 1 octet

proto_length - 1 octet

protocol (variable)

length of address in bytes - 1 octet

aggregator's address

For IP this encoding would be:

proto_type: 1 (ISO/IEC TR 9577 IPI/SPI)

proto_length: 1

protocol: hexadecimal 'CC'
length of address: 4
address: IP address

This attribute shall be handled as follows:

An IDRP for IP speaker may add this attribute to indicate that it performed aggregation.

5.2.6 SDRP Attribute

A new optional transitive attribute, SDRP is defined for use in IDRP for IP. This attribute is intended to both facilitate interoperation between IDRP for IP and BGP-4, and to allow SDRP speakers to be identified in the network. The type code of this attribute shall be 131. This attribute shall have the following encoding for each protocol family. Multiple protocol families may be included in the attribute.

proto_type - 1 octet)
proto_length (1 octet)
protocol (variable)
count of SDRP speakers - 2 octets
addresses of SDRP speakers
(The address information is specific to protocol.)

For IP Address, the format of the rest of the attribute is 4 byte octets, just like BGP.

For CLNP addresses, the format of the rest of the attribute is a set of tuples with (length of NETs, NET) with the following format.

length of address of SDRP speaker - 1 octet
address of SDRP speaker - (variable)

For IP this encoding would be:

proto_type: 1 (ISO/IEC TR 9577 IPI/SPI)

proto_length: 1

protocol: hexadecimal 'CC'

count of SDRP speakers

address: IP address list

This attribute shall be handled as follows:

An IDRP for IP speaker may add this attribute to if it is an SDRP speaker for the domain. It may add its own address and other SDRP speaker for the domain.

[5.2.7](#) Confederations of Autonomous Systems.

IDRP supports the ability to group Routing Domains into a Routing Domain Confederation. Likewise, IDRP for IP supports the ability to group a collection of Autonomous Systems into a Confederation of Autonomous Systems. With respect to the IDRP document in the context of IDRP for IP, the terms "Routing Domain Confederation" and "Confederation of Autonomous Systems" should be treated as synonymous.

[5.2.8](#) Domain Configuration Information

Correct Operation of IDRP described in [\[1\]](#) assumes that a minimum amount of information is available to both the inter-domain and intra-domain routing protocols. This information is static in nature, and is not expected to change frequently. This document assumes that this information is supplied via IDRP MIB ([\[6\]](#)). While the following is phrased in terms of MIB, this document allows alternative mechanisms (e.g. configuration files) as well.

The information required by a BIS that implements the IDRP for IP protocol is:

- a) Location and identity of adjacent Intra-Domain ISS (routers)

The MIB table INTRA-IS lists the IP addresses of the routers to which the local BIS may deliver an inbound NPDU whose destination lies within the BIS's routing domain. These routers listed in the INTRA-IS table support the intra-domain routing protocol of this autonomous system, and share at least one common subnet with the BIS.

In particular, if the local BIS participates in both the inter-domain routing protocol (IDRP) and the intra-domain routing protocol, then the IP address of the local BIS will be listed in the INTRA-IS table.

b) Location and identity of BISSs in the BIS's domain

This information permits a BIS to identify all other BISSs located within its routing domain. This information is contained in the MIB table INTERNAL-BIS, which contains a set of IP addresses which identify the BISSs in the domain.

c) Location and identity of BISSs in adjacent domains:

Each BIS needs information to identify the IP address of each BIS located in an adjacent RD and reachable via a single subnetwork hop. This information is contained in the IDRP MIB table EXTERNAL-BIS-NEIGHBORS, which is a table of IP addresses.

d) IP network address information for all systems in the routing domain

This information is used by the BIS to construct its network layer reachability information. This information is contained in the MIB table INTERNAL-SYSTEMS. The IP network address information shall be expressed in terms of IP address prefixes. A single prefix can be used to describe:

- a host address,
- a subnetwork number,
- a network number, or
- a collection of contiguous network numbers

e) LOCAL RDI

This information is contained in managed object LOCAL-RDI; it is the RDI of the routing domain in which the BIS is located. As specified in [section 7](#) of this document, the RDI for an IDRP for IP routing domain has an NSAP Address format. This NSAP Address format is built out of a fixed "header" and the autonomous system number of this autonomous system (routing domain).

f) RIB-AttSet

The RIB-AttSet contains information about the QoS functions a BIS will support. As described in [section 4](#), this can be none, any, or all of the Transit Delay, Residual Error, and Expense distinguishing attributes.

Expiration Date February 1994

[Page 11]

INTERNET DRAFT

September 1993

g) RDC-Config:

This information identifies all the routing domain confederations (RDCs) to which the RD of the local BIS belongs, and it describes the nesting relationships that are in force between them. It is contained in the MIB table RDC-Config.

Each RDC is identified by an RDI which has the format described in [section 7](#) of this document.

Note that since a domain is not required to belong to a confederation this information is optional and needs to be present only at BISs of the domains that are part of one or more of RDCs.

h) Local IP addresses

The LOCAL IP MIB table contains a list of IP addresses assigned to the interfaces of a BIS. This information is used to determine what interface should be used to forward outgoing NPDUs.

[5.2.9](#) Forwarding of IP packets

This section is intended to define the same function for IP packets as is defined for CLNP packets in the "Forwarding Process for CLNS" (Section 8 in [1]).

It is assumed that a BIS is capable of distinguishing between a FIB constructed by means of an intra-autonomous system routing protocol and a FIB constructed by means of IDRP.

After a BIS determines the packet's destination IP address, the BIS shall proceed as follows:

- a) If the destination address specifies a system located within the autonomous system of the receiving BIS, then the BIS shall forward the IP packet to any of the ISSs listed in the managed object INTRA-IS. That is, any further forwarding of the IP packet is the responsibility of the intra-autonomous system routing protocol; otherwise,
- b) the destination system is located in a different autonomous system, and the local BIS shall perform the following actions:

It shall determine the IP-Derived distinguishing attribute, according to clause 5.2.3. It shall next apply the procedures of clause 5.2.3 to determine if the IP-Derived distinguishing attribute matches any of the RIB-Atts of the information base(s) supported by the local BIS. If such a match is found, then the FIB with the matched RIB-Atts is used for forwarding.

If the procedures of clause 5.2.3 identify a non-default IP-Derived distinguishing attribute, but the local BIS does not maintain a FIB with the matching RIB-Atts, or the local BIS maintains a FIB with the matching RIB-Atts but this FIB does not have a route to the destination, then the local system sets the Type Of Service field in the IP packet to 0 and uses the FIB with no distinguishing attributes.

The incoming IP packet shall be forwarded based on the FIB entry that has the longest IP address prefix that matches the destination of the incoming IP packet, as follows:

- 1) If the entry in the inter-domain FIB that corresponds to

the destination address of an incoming IP packet contains a NEXT_HOP that identifies an interface of a BIS such that the interface is attached to a subnet shared with the local BIS, then the IP packet shall be forwarded directly to the BIS indicated in the NEXT_HOP entry over that interface; otherwise,

2) if the entry in the inter-domain FIB that corresponds to the destination address of the incoming IP packet contains a NEXT_HOP entry that identifies an interface of a BIS such that the interface is not attached to any of the subnets attached to the local BIS, then the local BIS has the following options:

i) Encapsulate the IP packet

The local BIS may encapsulate the IP packet, using its own IP address as the source address and the IP address of the next-hop BIS contained in the NEXT_HOP entry as the destination address. Since IP doesn't have a standard encapsulation protocol, use of this option should be discouraged.

ii) Use paths calculated by the intra-autonomous system routing protocols

The local BIS may query the FIB constructed by the intra-autonomous system routing protocols to ascertain if that FIB contains a route to the destination system. If that is the case, and if the path constructed by the intra-autonomous system routing protocols delivers the IP packet to the appropriate next-hop BIS, then the IP packet may be forwarded using the FIB constructed by the intra-autonomous system routing protocols.

IP_EXTFWD	Does the BIS correctly forward 5.3 IP packets with destinations outside its routing domain?	M	Yes___
IP_INTFWD	Does the BIS correctly forward 5.3 IP packets with destinations inside its routing domain? -----	M	Yes___

Table 1: PICS Proforma for IDRP: IP packet forwarding

The "ITEM" column describes the feature in the IP forwarding function that the IDRP implementation is to provide. The "Question/Feature" section seeks to describe the feature. The Reference is the section in this document that describes this feature. The status gives an indication of "M" - Mandatory feature for an IDRP implementation or "O" - optional feature. The "Support" column is a column for the implementor to check whether this feature is available in a particular implementation.)

[5.3](#) Advertising NLRI information for IP addresses

The NLRI field in an UPDATE PDU contains IP address information about systems that reside within a given routing domain or whose IP address space is under the control of the administrator of the routing domain. It should not be used to convey information about the operational status of these systems. The information in the NLRI field is intended to convey static administrative information rather than dynamic transient information; for example, it is not necessary to report that a given system has changed from offline to online.

End systems (hosts) and Intermediate systems (routers) within a RD using IDRP may use any IP address that is valid within the IP context. Within the NLRI, the address information for a set of IP addresses may be represented by an IP prefix. An IP prefix is the sequence of bits in a 4 byte IP address which are common between a set of IP addresses.

For example, the addresses 192.5.0.0 through 192.5.255.255 have the first 16 bits of the address information in common. These 16 bits of the IP address may be called an IP prefix which represents the set of

IP addresses 192.5.0.0 through 192.5.255.255.

Expiration Date February 1994

[Page 14]

INTERNET DRAFT

September 1993

An IP prefix must contain a contiguous set of bits starting at the most significant bit, but the bits may cover any part of the 4 byte IP address. The following guidelines for inclusion of IP address prefixes in the NLRI field of the UPDATE PDUs originated within a routing domain will provide efficient use of this protocol:

- a) Only IP prefixes representing IP addresses that are within the control of the administrator of a given routing domain may be reported in the NLRI field for a RD. These IP prefixes can represent IP addresses for systems which are:
 - online,
 - offline, or
 - allocated to the network, but not yet allocated to a machine.
- b) IP prefixes representing IP addresses outside of the RD administrator's control shall not be included in the NLRI.
- c) For efficient use of the protocol, the WITHDRAW ROUTES field should not be used to report the NLRI of systems that are offline. This field should be used only to advertise IP prefixes for IP addresses that are no longer under the control of the administrator of the local routing domain, regardless of whether the systems are online or offline.

A conformant implementation is required to have the ability to specify when an aggregated route may be generated out of partial routing information. A BIS at the border of an autonomous system (or group of autonomous systems) must be able to generate an aggregated route for a whole set of NLRIs over which it has administrative control, even when not all of them are reachable at the same time.

NEXT_HOP information associated with a particular route shall be derived from the NEXT_HOP attribute in the UPDATE BISPDU that carries the route. If that attribute is not present, it shall be derived from the source IP address of the IP packet that carries the UPDATE BISPDU containing the route.

7 Routing Domain Identifiers used for both IP and OSI

Routing Domain Identifiers (RDIs) are identifiers used in BISPDU's to uniquely identify individual routing domains and routing domain

Expiration Date February 1994

[Page 15]

INTERNET DRAFT

September 1993

confederations.

For ease of administration, the RDI is taken out of the NSAP address space. However, the RDI is just a number which identifies the routing domain, and need not bear any relationship to any reachable addresses within the domain.

For ease of interworking with other IP inter-autonomous system routing protocols, The RDI used for an autonomous system running IDRP for IP should be derived from an appropriately assigned Autonomous System Number as follows:

47:00:05:c0:01:aa:aa

Where 47:00:05:c0:01 is a unique prefix assigned by a valid addressing authority (NIST), and aa:aa is the autonomous system number.

This encoding of the RDI contains a "fixed header" (the 47:00:05:c0:01 sequence) plus the AS value.

(Note: While AS values are currently 2 octets long, IDRP allows Routing Domain Identifiers to be of arbitrary length. Thus, if the space of AS numbers is expanded to be greater than two octets, this may be accommodated by simply lengthening the above encoding--the AS number following the fixed header is considered to be right

justified, and its size can be unambiguously determined from the RDI length.)

[8](#) Deployment Guidelines for IDRP for IP

The correct and efficient operation of the IDRP protocol requires that certain guidelines are used for deployment with ISO routing Domains. Some equivalent deployment guidelines for IDRP for IP are required within Autonomous Systems. These guidelines represent only the required deployment guidelines, and not details on the usage of IDRP for IP in the Internet.

[8.1](#) Minimum configuration of an Autonomous System

An autonomous system using IDRP for IP must as a minimum contain:

one BIS, and one BIS capable of delivering NPDUs to the intra-domain routing function if the AS contains hosts.

Expiration Date February 1994

[Page 16]

INTERNET DRAFT

September 1993

[8.2](#) Multiple IDRP sessions between the same pair of routers

An IP router may have multiple IP addresses, one for each interface. In contrast, an OSI Intermediate System has only one Network Entity Title (network address). An OSI BIS thus may not have multiple IDRP sessions with another BIS, since the NET is unique and there is no mechanism for multiplexing sessions. However, an IP router may potentially have multiple IDRP sessions with another router, since each BIS may have multiple IP addresses, and one BIS may not be able to ascertain that those addresses correspond to the same BIS. Multiple IDRP sessions between IP BISs may not be efficient, but they are not illegal, nor do they impact the robustness of the IDRP for IP protocol; they will simply appear as multiple paths to the same neighboring AS. One possible way of avoiding multiple parallel IDRP sessions between a pair of BISs within a single autonomous system is

to bind all source addresses of outgoing BISPDUs to the IP address of a particular interface of the BIS. Likewise, for a pair of BISs located in adjacent autonomous systems, binding the source addresses to a single address of an interface attached to a common subnetwork provides for the elimination of multiple parallel sessions.

9. Recommended set of supported routing policies.

Policies are provided to IDRP in the form of configuration information. This information is not directly encoded in the protocol. Therefore, IDRP can provide support for very complex routing policies (an example of such policy is presented in Annex K of [1]). However, it is not required that all IDRP implementations support such policies.

We are not attempting to standardize the routing policies that must be supported in every IDRP implementation, but we strongly encourage all implementors to support the following set of routing policies:

IDRP implementations should allow an AS to control announcements of IDRP-learned routes to adjacent AS's. Implementations should also support such control with at least the granularity of a single network. Implementations should also support such control with the granularity of an autonomous system, where the autonomous system may be either the autonomous system that originated the route, or the autonomous system that advertised the route to the local system (adjacent autonomous system). Care must be taken when a BIS selects a new route that can't be announced to a particular external peer, while the previously selected route was announced to that peer. Specifically, the local system must explicitly indicate to the peer that the previous route is now infeasible. IDRP implementations should allow an AS to prefer a particular path to a destination when more than one path is available. This function may be implemented by allowing system administrators to assign "weights" to AS's and having the route selection process select a route with the lowest "weight" (where "weight" of a route is defined as a sum of "weights" of all AS's in the RD_PATH path attribute associated with that route). IDRP

implementations should allow an AS to ignore routes with certain AS's in the RD_PATH path attribute. Such function can be implemented by using the technique outlined in [9], and by assigning "infinity" as "weights" for such AS's. The route selection process must ignore

routes that have "weight" equal to "infinity".

10. Security Considerations

Security issues are not discussed in this document.

11. Acknowledgements

A large set of thanks to Dave Katz (cisco) who helped edit the document. We would also like to express my thanks to Scott Brim (Cornell University) for his review and constructive comments.

The authors would like to acknowledge contributions made by authors of [8], Yakov Rekhter and Phill Gross. Certain sections of this document are taken (sometimes almost verbatim) from their document.

12. References

[1] ISO/IEC IS 10747 - Information Processing Systems - Telecommunications and Information Exchange between Systems - Protocol for Exchange of Inter-domain Routeing Information among Intermediate Systems to Support Forwarding of ISO 8473 PDUs, 1993.

[2] ISO 8473 - Information Processing Systems - Data Communications - Protocol for Providing the Connectionless-mode Network Service, 1988.

[3] ISO/IEC 10589 - Information Processing Systems - Telecommunications and Information Exchange between systems - Intermediate System to Intermediate System Intra-Domain routeing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473), 1992.

[4] ISO 9542 - Information Processing Systems - Telecommunications and information exchange between systems - End system to Intermediate system routeing exchange protocol for use in conjunction with the Protocol for providing the connectionless-mode network service (ISO 8473)

[5] [RFC 791](#) (Jon Postel, editor) - Internet Protocol - DARPA Internet Program Protocol Specification (September 1981)

[6] Hares, S., "Definition of Managed Objects for the IDRP for IP",

INTERNET DRAFT

September 1993

Internet Draft, March 1993

[7] Fuller, V., Li, T., Yu, J, and Varadhan, K., "Supernetting: an Address Assignment and Aggregation Strategy", [RFC1519](#), September 1993

[8] Rekhter, Y., Gross, P., "Application of the Border Gateway Protocol in the Internet", Internet Draft September 1992

[9] Braun, H-W., "Models of Policy Based Routing", [RFC 1104](#), Merit/NSFNET, June 1989.

[10] Rekhter, Y., Li, T., "An Architecture for IP Address Allocation with CIDR", [RFC1518](#), September 1993

[11] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4), Internet Draft, cisco Systems, T.J. Watson Research Center, IBM Corp., September 1992.

[12] Almquist, P., "Type of Service Routing in the Internet Protocol Suite", [RFC 1248](#), July 1992.

Authors' Addresses

Susan Hares
Merit, Inc
1071 Beal Avenue
Ann Arbor, MI 4810x

Phone: (313) 936-2095
Email: skh@merit.edu

John Scudder
Merit, Inc
1071 Beal Avenue
Ann Arbor, MI 4810x

Phone: (313) 764-5384
Email: jgs@merit.edu

IETF IDRP for IP WG mailing list: idrp-for-ip@merit.edu
To be added: idrp-request@merit.edu

Expiration Date February 1994

[Page 19]