    New Scheme for Internet Routing and Addressing (ENCAPS) for IPNG

Status of This Memo

Abstract

   This document was submitted to the IETF IPng area in response to RFC
   1550. Publication of this document does not imply acceptance by the
   IPng area of any ideas expressed within. Comments should be
   submitted to the big-internet@munnari.oz.au mailing list.

   This memo describes a proposal made to to the Routing and Addressing
   group [ROAD] January 1992 by Robert Hinden. It was originally sent
   as an email message. It proposes a medium term solution to the
   Internet's routing and addressing problems.

INTRODUCTION

   I would like to propose a new scheme which I believe is a good medium
   term solution to the routing and address problems of the internet.
   It has the following positive attributes:

      - No Changes to Hosts
      - No Changes to Most Routers
      - No New Routing Protocols
      - No New Internet Protocols
      - No Translation of Addresses in Packets
      - Reduces the Routing Table Size in All Routers
      - Uses the Current Internet Address Structure

   It is not a solution good for all time, because it does impose some
   size limits and does not support new internet services such as
   guaranteed bandwidth, delay, etc. It does require border routers to
   do additional processing, but does not require any packet
   translation. I believe that this scheme will give us enough time to
   put into place a long term solution (i.e. pick one or more of CLNP,
   *NAT, IDPR, IDRP, Nimrod, Unified, NewIP, etc.)

This scheme is based on the ideas presented by Deborah Estrin (route
on ADs), Martha Steenstrup (encapsulation), and probably steals from
ideas put forward by Noel Chiappa, Van Jacobson , Ross Callon, Dave
Oran, and everyone else in the ROAD group.

CONTEXT

I think that we (the ROAD group) agree that in the short term we need
to make better use of the IP address space. I think we also (mostly)
agree that in the long term we need a solution that can deal with a
very large number of end points and routes, as well as support new
services such as guarantees of service, source selected routes, etc.
We do not agree on any of the details of this but do agree that we
can not figure out a long term solution before March. We do agree
that we should start working on a long term solution(s).

What this leaves is the need for a good medium term solution which
can keep the Internet going until we can design and deploy a long
term solution. The medium term solution wants to be the most "cost
effective". It should buy us the most time to develop a long term
solution and do it with as little change to the existing Internet as
possible.

I propose this scheme as a new medium term solution.

NEW SCHEME

The basic idea is that inter-domain routing be done by routing on
autonomous domains (AD). The key is how this is done. The mechanism
to do this is for the border routers to encapsulate the original IP
datagrams with another IP header. The source and destination
addresses in the new header (I will call it the AD-Header from here
on) represent the source and destination ADs.

When the first (entrance) border router receives a datagram from a
host or router without an AD-Header it looks at the source and
destination address and does a DNS lookup to get the addresses for
the AD-Header. It then adds an AD-Header and forwards the
encapsulated datagram to its proper destination AD.

The border routers would compute AD routes by running a routing
protocol between themselves. BGP or even IS-IS or OSPF for that
matter, would work fine. As you will see later, they might even be
better.

The addresses I propose to use for the AD addresses are plain old IP
addresses. A small number of Class A and Class B addresses would be
reserved for this purpose. The network number of the address would

indicate that it was an AD identifier. The local part of the address
would indicate the actual AD. This would allow for many ADs to be
supported. For example, 10 Class-A and 10 Class-B addresses could
accommodate (10*2^24 + 10*2^16) 168,427,500 ADs. We clearly don't
need that many for a long time.

The reason why I would chose to get more than one network number to
use to represent the AD address is I would use them to organize the
ADs. Let's call them commonwealths. Each commonwealth would only
have to know the detail of it's own ADs.

Next I would have the border routers inject these AD addresses into
the Intra-AD routing of transit ADs. They would tell the routers
inside of the transit AD that they (the border routers) were the
route to each appropriate AD network. Commonwealths that have
multiple interconnects (probably the common case) could by the use of
careful assignment of the AD addresses use subnetting to support
reasonable routing between the commonwealths. This is where OSPF or
IS-IS might be better than BGP. Also, IS-IS, with its ability to
route on actual end points might be the best.

The motivation behind injecting the AD addresses into the Intra-AD
routing of the transit ADs, is that the routers in these ADs can
forward the AD-Headers without knowing that they are special. Only
the entrance and exit border routers are required to do anything
different.

Finally when a AD-Header is received at the last (exit) border router
it strips off the AD-Header and sends the datagram to the final
destination.

This scheme is based around the idea that IP addresses are globally
unique. I think that we will not actually run out of IP addresses
for a long time and that we can live with the current addressing
until we can deploy a long term solution.

This scheme could be extended to not require globally unique IP
address. Effectively the combination of AD-Address and IP-Address is
the globally unique address. To use this scheme without globally
unique IP-Addresses and without changing in the hosts would require a
NAT mechanism in the border routers. I think it would be preferable
to change the hosts to have them do the DNS query and add the AD-
header. This could be the basis for the long term solution.

Another interesting aspect of this scheme is that if we were to relax
the current architecture where one IP-Address is always in only one
AD, to allow an IP-Address to be in more than one AD, it would
provide a solution to the issue of allowing a IP entity to get

service from more than one service provider.

SUMMARY OF CHANGES REQUIRED

The DNS needs to be extended to add an AD-Address entry for each
name. These will be used by the entry and exit border routers to get
the AD-Addresses to use when building the AD-Headers.

Border routers need to be extended to do the DNS lookup, perform AD-
Header encapsulation, run an inter-AD routing algorithm using AD-
Addresses, and be able to AD-Header de-encapsulation.

CONCLUSION

I believe that this scheme has may advantages. These are:

   - Only border routers and the DNS need change. No changes are
     required in hosts or non-border routers.

   - No performance impact on datagram forwarding except at entry
     and exit border routers.

   - Only a small impact on bandwidth utilization on transit
     networks due the addition of a 20 byte IP header to each
     datagram.

   - Removes the Inter-AD routing from Intra-AD routing and as a
     result solves the routing load (table size and computation)
     problem for the foreseeable future.

   - The routing load on the border routers is manageable because
     border routers only need to know the detail of the routing
     commonwealth they are a member of. Other commonwealths appear
     as single addresses.

   - No requirement for new routing protocols to be designed or
     deployed.

   - No translation of packets from one address scheme to another.

   - Uses the current IP addressing structure.

   - It scales well even if there is on the order of one AD per IP
     network, because the AD-Addresses can be assigned logically.

It does have some disadvantages. These are (at least):

    - It is not a long term solution in its initial form.

    - It assumes that the current IP-Addresses can remain globally
      unique for a long time.

REFERENCES

    [ROAD] Gross, P., and P. Almquist, "IESG Deliberations on Routing
           and Addressing", RFC 1380, ANS, Stanford University,
           November 1992.

SECURITY CONSIDERATIONS

    Security issues are not discussed in this memo.

AUTHOR'S ADDRESS

    Robert M. Hinden
    Ipsilon Networks, Inc.
    2191 East Bayshore Road
    Suite 100
    Palo Alto, CA 94303
    USA

    EMail: hinden@ipsilon.com
    Phone: +1 (415) 846-4604