                  Guidelines for Rate Limits on the MBONE

1. Status of this Memo

   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
   and may be updated, replaced, or obsoleted by other documents at any
   time. It is inappropriate to use Internet-Drafts as reference
   material or to cite them other than as ``work in progress.''

   To learn the current status of any Internet-Draft, please check the
   ``1id-abstracts.txt'' listing contained in the Internet-Drafts Shadow
   Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe),
   munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or
   ftp.isi.edu (US West Coast).

2. Abstract

   Much confusion and misunderstanding exists in the Internet community
   on the use of rate limits on the Multicast Backbone (MBONE). This
   document dispels some of the misunderstandings of rate limits and
   describes the best current practices for when to implement rate
   limits on MBONE links.

   It is a product of the Multicast Deployment Working Group in the
   Operational Requirements area of the Internet Engineering Task Force.
   Submit comments to <mboned@ns.uoregon.edu> or the author.

3. History of DVMRP Tunnel Rate Limits

   The MBone (Multicast Backbone) of the Internet is composed of a DVMRP
   backbone connected to regions that may be running other multicast
   routing protocols. These multicast islands are connected together by
   DVMRP tunnels across the unicast Internet backbone.

   The original multicast router for the MBONE was mrouted. Mrouted,
   through version 3.8, has a rate limit of 512 kbps enabled by default

on each tunnel interface. This rate limit was included in mrouted
because of concern that multicast traffic could consume all available
bandwidth on some links of the Internet. Other multicast router

implementations include support for rate limits on interfaces but
they are not enabled by default.

Several things have changed in the Internet and IP multicasting that
make the idea of a global rate limit for the MBONE no longer
necessary. First, the size of provider's backbone circuits has
increased greatly. It is no longer a valid concern that a single
multicast transmitter could send enough multicast traffic to
completely fill a national providers backbone circuit since those
circuits have moved from DS-1 (1.536 Mbps) to DS-3 (44.210 Mbps).

Second, multicast pruning is now widely deployed across the MBONE and
is being mandated as a requirement for MBONE multicast routers
[PRUNING]. Pruning prevents multicast routers from forwarding
multicast traffic to neighbors that have no receivers for the
destination multicast group. With the original implementation of
mrouted, pruning was not implemented so multicast traffic sent to the
MBONE was distributed across the entire backbone.

Because of the original lack of pruning and the default rate limiting
of 512 kbps enabled on mrouted, many people believe that there is a
total capacity of 512 kbps for the aggregate traffic sent on the
MBONE. In fact, many people have used this as a reason not to
develop multicast applications. With increased bandwidth in the
national service providers networks and with pruning multicast
routers deployed widely on the MBONE, neither of the previous
constraints on MBONE traffic are valid any longer.

4. Current Use of Rate Limits

Backbone providers that are tunneling multicast traffic across their
networks are encouraged to raise the default rate limit on mrouted
boxes to a more useful limit. It is essential to encourage the use
of multicast traffic as a more scalable alternative to unicast
streaming technologies.

This is not to say that rate limits no longer are required anywhere
on the MBONE. End-points on the MBONE may still require rate limits

their tunnels to prevent multicast users from consuming all available
bandwidth on the customer's access link. End-points should determine
how much of their bandwidth they want to allow multicast traffic to
consume and should request that their tunnel provider only allow that
much traffic traffic on the tunnel. Some multicast routers also
permit the configuration of an inbound rate limit to prevent
congestion within the router.

Rate limits are also very useful on congested links due to the fact
that the current DVMRP implementation does not retransmit prunes. If
the prune message is lost, the group will continue to be transmitted
across the tunnel.

Backbone providers may also choose to set a rate limit on their
tunnels that is high enough not to impede normal multicast traffic
but low enough to protect against denial-of-service attacks. The
potential for denial of service attacks on a multicast group are
similar to "flood pinging" against a unicast targe, except the
potential exists to hit more that one path through the network at the
same time.

In all applications of rate limits, it is important to remember that,
with pruning, these rate limits will only apply to multicast groups
that have receivers at the other end of the tunnel. At no time does
a rate limit applied to a DVMRP tunnel imply a total aggregate
capacity of the MBONE.

5. Security Considerations

The removal of rate limits altogether on backbone provider's MBONE
tunnels could possibly lead to denial-of-service attacks by sending
large amounts of traffic to a global multicast group such as the SDR
group.

6. References

[PRUNING] J. Hawkinson, "Multicast pruning a necessity", Work in
        progress (internet-draft).

7. Author's Address

Doug Junkins
NorthWestNet
15400 30th Place SE, Suite 202
Bellevue, WA 98007

phone: +1 206 649 7419
email: junkins@nwnet.net