The Multicast Backbone (MBone) Network

ABBREVIATION

AS: Autonomous System

BOOTP: Bootstrap Protocol

CBT: Core-Based Tree

DNS: Domain Name System

IDMR: Inter-Domain Multicast Routing

IETF: Internet Engineering Task Force

IGMP: Internet Group Management Protocol

DVMRP: Distance Vector Multicast Routing Protocol

MOSPF: Multicast extensions to Open Shortest Path First

OSPF: Open Shortest Path First

PIM: Protocol Independent Multicast

PIM-DM: Protocol Independent Multicast - Dense Mode

PIM-SM: Protocol Independent Multicast - Sparse Mode

RIP: Routing Information Protocol

RPB: Reverse Path Broadcast

RPM: Reverse Path Multicast

ST: Steiner Tree

TRPB: Truncated Reverse Path Broadcast

TTL: time-to-live

MBone: Multicast Backbone

1. INTRODUCTION

Today's Internet faces a series of serious challenges that were largely nonexistent as its inception. Not only is the number of users of the Internet growing exponentially, but so too are the number of networks comprising the Internet, and thus the number of devices used to interconnect these networks. This unprecedented growth is accompanied by the multiplication of new applications. Many applications previously available only to limited numbers of power users with high-end workstations are starting to become mainstream applications in the PC world – videoconferencing, video broadcasts, for example. All these new applications bring with them new types of data – video and audio, for example – which make new demands on a network's response and ability to deliver the data. Network designers are now faced with the challenge of supporting the timely and reliable delivery of any kind of data, especially real-time and multimedia data, to any user.

In addition to the new requirements for real-time data, applications are quickly evolving from one-to-one communications to one-to-many and many-to-many communications. Widespread use of these applications can easily overload existing networks when the same bits of information have to be transmitted to different users at the same time. But new technologies using more intelligence distributed across the network make it possible to reduce unnecessary duplication of bites and relieve some of the network load. This article will focus on one of the new technologies that's rapidly growing in use IP multicasting.

IP multicasting owes its heritage to an experimental service on the Internet called the MBone, or Multicast Backbone. A virtual network layered atop sections of the physical Internet, the MBone was designed to create a semipermanent IP Multicast testbed without waiting for the deployment of multicast-capable routers throughout the entire Internet. Although the MBone is an important resource and testbed, it has some limits in performance and scalability that will keep it from becoming a truly enterprise-class multicast service. This article will describe the architecture, limitation and future development of MBone.

2. BACKGROUND OF MBONE

The MBone had its origins in a joint research project of the University of Southern California's Information Sciences Institute, the Massachusetts Institute of Technology, the Xerox Palo Alto Research Center, the Lawrence Berkeley National Laboratory, and other sites sponsored by the U.S. Defense Advanced Research Projects Agency. In 1990 this community created a wide-area research network called the DARPA Research Testbed Network, or DARTNet, consisting of Unix workstations serving as programmable routers interconnected via T1 links. Early in the project, the DARTNet community deployed a preliminary version of IP Multicast over the network. This provided the first opportunity to experiment with wide-area network-layer multicast on a nontrivial scale.

The next major step occurred when the DARTNet community developed several real-time, interactive multimedia applications using IP Multicast as the network layer to study the problem of multiparty remote conferencing over packet-switched networks. The utility of the weekly DARTNet meetings generated interest in extending the multicast infrastructure beyond the reaches of the testbed. At that time, however, production Internet routers could not carry multicast traffic. The designers of IP Multicast therefore enabled multicast routers to forward packets and exchange routing messages over virtual links.

The ability to build multicast subnets using tunnels prompted an ambitious experiment in March 1992 (see Figure 1). Thirty-two isolated multicast sites spread over four countries were configured into a large, virtual multicast network, which in turn was used to audiocast the 23rd Internet Engineering Task Force meeting. The virtual network backbone used to glue together the multicast-capable subnetworks was called the “Multicast Backone”, or MBone. Despite a few technical difficulties, the experiment proved enormously successful. In March 1997, it included 3,400 multicast-enabled subnets tied together. In just a few years, the MBone evolved from a small research curiosity to a large-scale, widely used communications infrastructure.

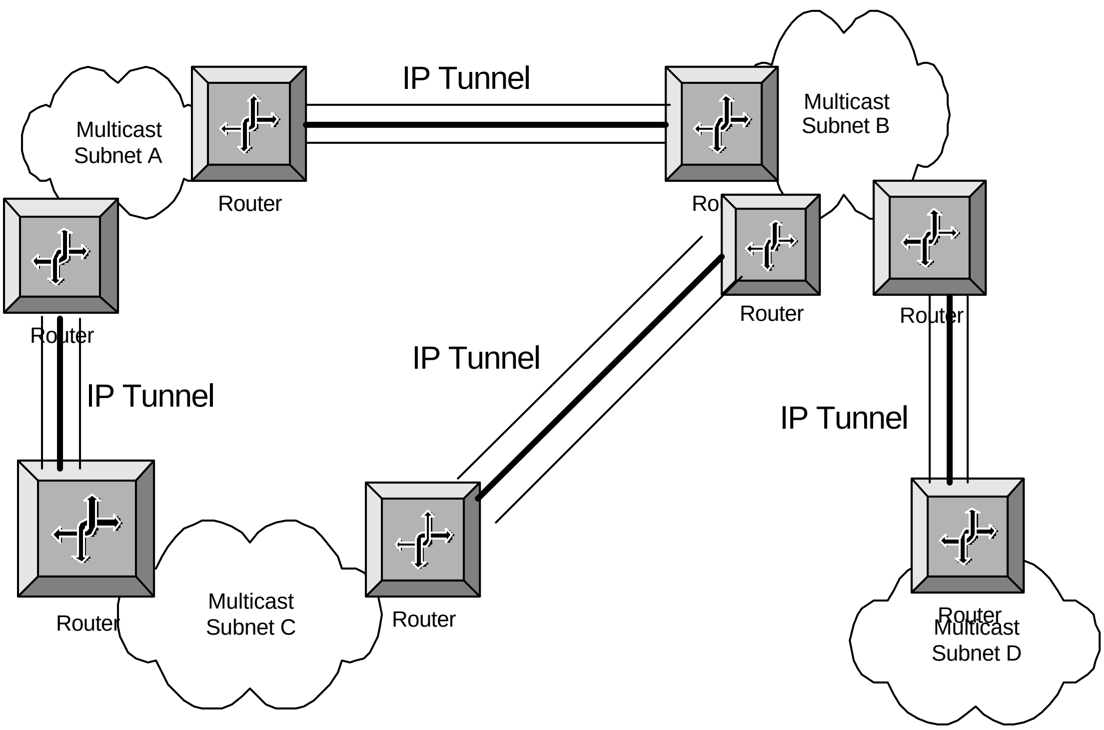

Figure 1 – Architecture of MBone

The next section will describe how IP Multicasting works.

3. IP Multicasting

From a networking standpoint, the multicasting single send operation that results in copies of the sent data being delivered to many receivers – can be implemented in two ways (see Figure 2):

(a) Sender will use a separate unicast transport connection to each of the receivers. An application-level data unit that is passed to the transport layaer is then duplicated at the sender and transmitted over each of the individual connections. The multicast sender uses separate unicast connections to reach the receivers.

(b) Sender transmits a single datagram. This datagarm is then replicated at a network router whenever it must be forwarded on multiple outgoing links in order to reach the receivers. This approach toward multicast makes more efficient use of network bandwidth in that only a single copy of a datagram will ever traverse a link. This approach is implemented in the Internet and utilises IGMP to identify and address all of the receivers.

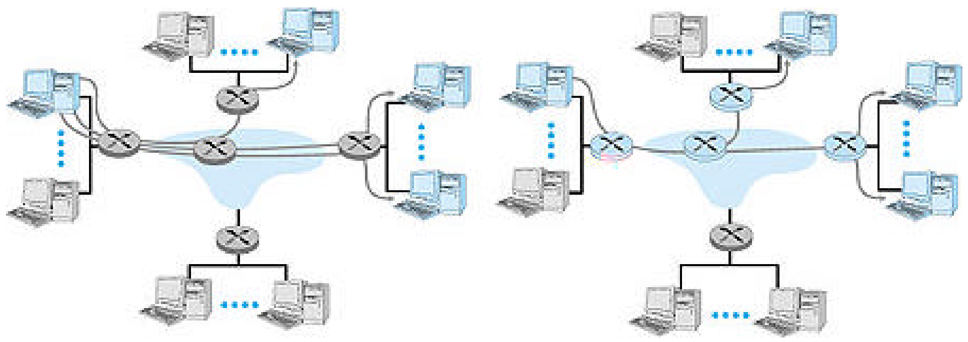

Figure 2 – Two approaches toward implementing the multicast abstraction

IP Multicast is already changing the ways in which companies do business. Now with support of IP Multicast across the Internet, it enables traditional business applications, such as videoconferencing, to scale across an enterprise's global locations. In addition, IP Mutlicast is being implemented to support e-commerce and integrated information.

4. <u>LIMITATIONS AND FUTURE OF MBONE</u>

As the MBone has grown, it has suffered from an increasing number of problems, and these problems have been occurring with increasing frequency.

a) Scalability - The most important reason for this is the growing difficulty of managing a flat virtual topology. As the MBone has grown, its size has become a problem, in terms of both routing state and susceptibility to misconfigurations. Large, flat networks are inherently unstable. Exacerbating this problem are organizational mechanisms which do not provide significant route aggregation. At its peak, the MBone had almost 10,000 routers.

b) Manageability – As the MBone has grown randomly, it has become harder to manage. The MBone has no central management, and most tasks have been handled on a per-site basis. Most coordination takes place via the MBone mailing list. Two types of inefficiency commonly observed are:

- Virtual topology (tunnel) management. The MBone is characterized as a set of multicast-capable islands connected by tunnels. The goal has always been to connect these islands in the most efficient manner, but over time suboptimal tunnels have been created. Tunnels are often set up in very inefficiency ways.
- Interfomain policy management. Domain boundaries are another source of problems when trying to manage a flat topology. The model in today's Internet is to establish autonomous system (AS) boundaries between Internet domains. ASes are commonly managed or owned by different organizations. Entities in one AS are typically not trusted by entities in another AS. As a result, exchange of routing information across AS boundaries is handled very carefully.

The first of above problems is the complexity and instability of a large flat topology. The second problem is that there are no protocol mechanisms to build a hierachical multicast routing topology. The need to solve these two problems created the first attempts to deploy interdomain maulticast.

Interdomain multicast has evolved out of the need to provide scalable, hierchical, Internet-wide multicast. Protocols that provide the necessary functionality have been developed, but the technology is relatively immature. These protocols are being considered by the IETF, while simultaneously being evaluated through extensive deployment.

Related WWW Links:

1. "Introduction to IP Multicast Routing,"
   http://www.3com.com/nsc/501303.html

2. "Unicast and Multicast Addressing and Routing Introductions,"
   http://www.ipmulticast.com/community/links-routing.html

3. "The IP Multicast Initiative,"
   http://www.ipmulticast.com/

4. "Hot Topics in Networking References,"
   http://www.cis.ohio-state.edu/~jain/refs/all_refs.htm

5. "The Internet Engineering Task Force,"
   http://www.ietf.cnri.reston.va.us/home.html

6. "A Tutorial on IP Multicast,"
   http://ganges.cs.tcd.ie/4ba2/multicast/antony/index.html

7. "MBone Resources,"
   http://www.mbone.com/mbone/references.html

8. "The MBone FAQ,"
   http://www.mbone.com/mbone/mbone.faq.html.

9. "Reliable Multicast Protocols,"
   http://www.tascnets.com/mist/doc/mcpCompare.html