

Transport of NetNews via IP-Multicast

Heiko W.Rupp
<hwr@pilhuhn.de>
18.02.1997

Abstract

The Bandwidth consumed by NetNews is growing steadily. While some time ago the rule of thumb was that it will double every year it seems now less dramatic, but still increasing. Even if NetNews used local caches since it starting (news spool at the provider) and a good flood fill algorithm for distribution, the bandwidth use of the physical links is still not optimal. Article transport via IP-Multicast can be a good way to improve link usage.

What is now

On the Internet distribution of NetNews is accomplished on a store and forward base. A host that gets an article, stores it locally in its databases and then offers it via NNTP[KL86] to all its configured neighbours that might be interested in the article. Neighbours that already have the article refuse to accept them so that they won't get transferred to a remote twice. This is a efficient way to distribute the articles, but it doesn't take into account that for offering and especially transfer of articles these often traverse the same physical link multiples times and thus imposing there a load that is much bigger than it has to be.

One way to oppose this load is to install servers on all nodes where backbone links meet and just let these servers talk to each other. The disadvantage of this system is that additional servers are needed which cost a lot of money.

What can be

Nowadays most backbone routers are already able to distribute multicast data, so one can use this infrastructure to get a mesh of distribution where articles don't cross links multiple times. If a site wants to get some newsgroups it just joins the appropriate multicast group. Most modern operating systems follow RFC 1112 [Dee89] and so have the prerequisite to send and receive news via multicast.

Kurt Lidl, Josh Osborne and Joseph Malcolm already described such a mechanism of distribution in [LOM94], but this one never got into production mode as it was too slow; also code is not free. There were also no provisions for spooling of articles if the remote server is down nor for compression of the articles before transmission.

There are basically two approaches to get NetNews over IP-multicast working:

Reliable multicast This one provides a way of distribution that, like TCP, guarantees you (within some bounds)

that the data you are sending off arrives at the receivers completely. For this to work you need either some additional hard- or software that provides the fault tolerance[Hof94, SDW92]. Or, if you let all receivers send their acknowledge packets directly to the multicast sender, you will get network and performance problems there (known as the implosion problem) [LS95].

Unreliable multicast Here IP packets with some payload are just put on the wire and flow to the receiver. Packets may get lost as depending on line quality and may arrive in wrong order. As packet loss rate is monitored in daily use and kept on a low level, this is no problem in practice¹.

The way to go is with unreliable multicast for the following reasons:

- Wrong article order doesn't matter as it is not important for news articles to arrive in order. Reader software takes care of that and different users want different sorting of articles anyway.
- There is no need for additional hard- or software which reduces costs of machines and administration.
- It doesn't matter much if some articles get lost as standard NNTP is a good way to deliver this lost articles reliably.
- There is no implosion problem.
- With most articles fitting in just one UDP packet (see below) the overhead of sending acknowledges and or keeping local copies (in the routers or dedicated hops) until all articles are reliably delivered is too high.
- Can be used on the existing MBONE [Dee93].

¹if packet loss happens due to link overload then a multicast feed instead of several unicast feeds can even improve the situation

After all this considerations, the following is done at the sender: For each article that is appropriate, compress it if you want, put it into a UDP packet, add some checksum and send it off the net. After some time offer the article to the remote via NNTP. At the receiver the packet is decoded and offered to the (local) news system via NNTP. If the local server is not available spool the article to a local disk. Article compression at the sender is done with the freely available Zlib [LD96].

Some data

Research has shown that over 90% of all articles are less than 64 kBytes in size² so that they fit well into the 64 kByte limit of UDP packets.

With a very alpha version of the code it was possible to get between 50 and 80 % of the articles via multicast even if a real time feed with *nntplink* was offering the same articles in parallel. If this NNTP feed is delayed for some time (e.g. 30 seconds) then the count of successfully transferred articles via multicast will even increase.

Conclusion

Transport of NetNews via multicast is a good way to go. If it is done "right", then it can even compete with the speed of fast classical transport mechanisms as *nntplink* while still saving bandwidth.

About the author

I am a undergraduate student of computer science at University of Karlsruhe, Germany. *Transport of NetNews via IP-Multicast* is my diploma thesis. To earn my cookies I am working at German ISP NTG/Xlink³ as administrator of Netnews.

The address for contact is:

Heiko W.Rupp
Gerwigstr.5

D-76131 Karlsruhe
Germany
Tel.: +49 721 966 1524
Fax: +49 721 661937
Email: hwr@pilhuhn.de

Sponsors are welcome!

References

References

- [Dee89] Steve Deering. *RFC 1112 – Host extensions for IP multicasting.*, August 1989. replaces RFC 1054 and RFC 988.
- [Dee93] S. Deering. *MBone: The Multicast Backbone*, M'arz 1993. CERFnet Seminar.
- [Hof94] M. Hofmann. *Zuverl'assige Kommunikation in heterogenen Netzen*, August 1994. Diplomarbeit am Institut f'ur Telematik, Universit'at Karlsruhe.
- [KL86] Brian Kantor and P. Lapsley. *Network News Transfer Protocol.*, Februar 1986.
- [LD96] J. Gailly L. Deutsch. *ZLIB Compressed Data Format Specification version 3.3*, Mai 1996.
- [LOM94] Kurt Lidl, Josh Osborne, and Joseph Malcolm. *Drinking from the firehose: Multicast usenet news.* Technical report, UUNET Technologies Inc., Januar 1994.
- [LS95] John C. Lin and Paul Sanjoy. *RMTP: A Reliable Multicast Transport Protocol*, Februar 1995.
- [SDW92] W. T. Strayer, D.J. Dmepsey, and B.C. Weaver. *XTP: The Xpress Transfer Protocol.* Addison-Wesley, Reading, Mass., USA, 1992.

²<http://www.xlink.net/hwr/histo/>

³<http://www.xlink.net/>