Not Dead Yet:
Viability of Multicast for Content Delivery

Sergey Alexandrov
Rutgers University
serge1@paul.rutgers.edu

Abstract

Multicast has always been designed with content delivery in mind, providing an efficient distribution transport. Expectations were for it to become the standard. However, today it is still not ready for prime time, with limited penetration, with few end users being able to access multicast content. Meanwhile, peer to peer technologies have skyrocketed, becoming the popular method for content. Some doom multicast to death by atrophy. In this paper, we aim to show that its demise is premature; rather, the advent of high-bandwidth end-system access points may precipitate the true rise of multicast, pushing it to the critical mass en route to a more global deployment.

1. Introduction

With the Internet constantly becoming a richer source and carrier of content, the question of content delivery becomes more and more important. No longer is simple one-to-one communication sufficient, since, as with most content, there are more consumers than there are providers; this inequality results in increased traffic and congestion.

Multicast was the part of the IP specification that appeared to be particularly fitting for content delivery because it allowed efficient one-to-many communication. It was envisioned that multicast's ability to efficiently distribute data to multiple targets would make it the primary method of delivery – many in the research and industry communities predicted multicast to be ubiquitous by this time. However, this has not come true – multicast exists in isolated network islands, and its use on the Internet remains limited.

In the meantime, other technologies have evolved to address multi-target communication needs, the most obvious and successful of these being P2P, peer-to-peer systems. Some are inclined to dismiss multicast as a dead technology. However, this is not necessarily the case. To determine this, we will examine multicast, to determine what it offers and why it has not yet succeeded, and compare it to the success of P2P approaches.

2.1 IP Multicast Overview – the Basics

An IETF (Internet Engineering Task Force) standard, IP multicast enables many-to-many communication between end hosts. It is based on a subscribe group membership model, where a multicast address designates a multicast group, and end hosts notify their routers of memberships in a particular group. Routers use this information to decide whether to forward a multicast packet, and whether it should be replicated to other interfaces.

IP multicast results in significant bandwidth savings. When delivering unicast content, a provider must send out a packet for each of the subscribers, resulting in a linear bandwidth cost across the entire path. When multicast, the provider needs only to generate a single packet, which will be replicated as necessary by routers along the paths, resulting in a tree-like routing path, where the "width" of any branch is only a single packet.

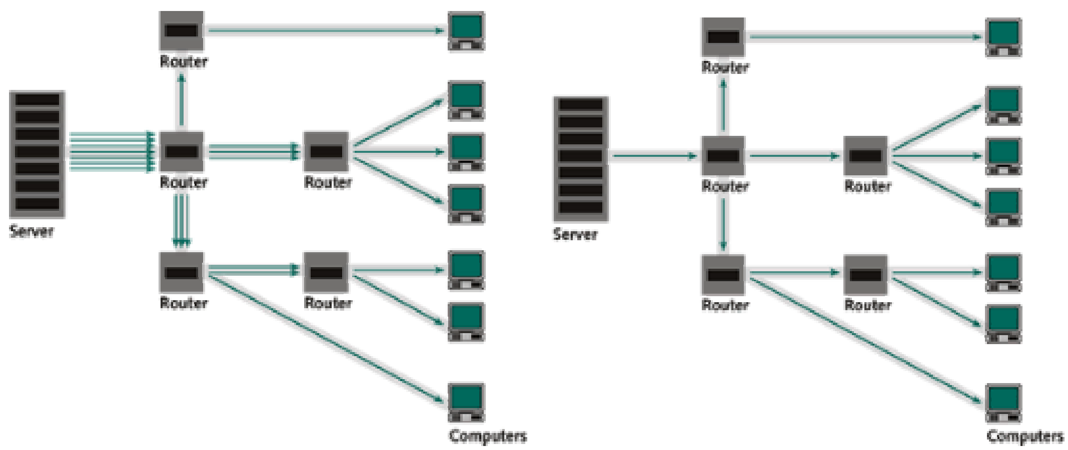

Figure 1: Unicast (left) vs. multicast (right)

For IP multicast to work, unicast methodologies alone are not sufficient. Multicast functionality requires a couple of parts:

Membership Management – router need to know who belongs to what groups in order to correctly forward packets. This was done by IGMP, Internet Group Management Protocol, which dictates how hosts declare their memberships, and how routers maintain it.

Routing – routers must make decisions where to route multicast packets, since multicast addresses function differently from unicast addresses. The first standard for this was DVMRP, Distance Vector Multicast Routing Protocol. This addressed intra-domain routing only, inter-domain routing was not yet properly addressed, initially implemented, or rather bypassed, via unicast tunneling.

The basic premise of multicast leads to a large set of potential applications, such as:

- Multimedia broadcasting (audio/video) – particular multi-channel/multi-camera options
- Teleconferencing/Videoconferencing
- Real-time collaborative applications (online instruction, for example)
- Multi-player network games
- Data replication
- Announcements, software updates
- Distributed computing

Multicast's benefits over unicast are obvious, and grow larger along with the number of receivers. This relationship is shown graphically in Figure 2.

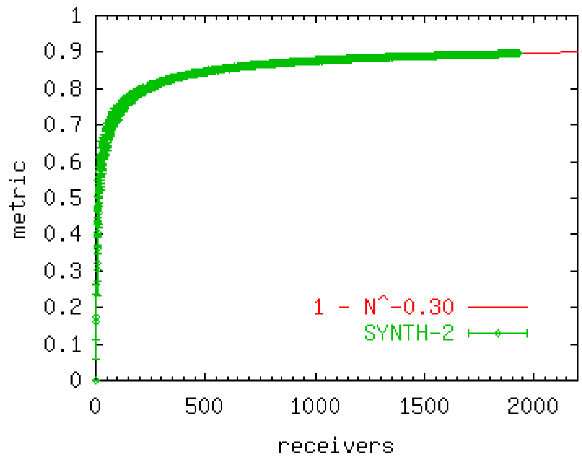

Figure 2: Percentage gain in bandwidth utilization achieved by using multicast rather than unicast (red is formulized prediction, green is observed data). [15]

2.2 Evolution of Multicast

Unfortunately, while the multicast basics are clear, implementations were not. For this reason, multicast initially was limited to research-oriented network islands, in particular MBONE, Multicast Internet Backbone. This is where much of the multicast research in the 1990s occurred, having to tackle several issues.

First, a better routing protocol was needed. DVMRP used reverse path forwarding type flooding to determine which routers were active for particular groups. This was costly, and would not scale in the greater Internet. The result of this was the PIM SM routing protocol – Protocol Independent Multicast, Sparse Mode. PIM SM forms source based distribution trees, and only when traffic is requested. A more efficient routing implementation, this is the current routing protocol standard for multicast.

Second, membership management models needed to be refined. The early model was ASM – any source multicast, where a subscriber received all content from a group, regardless of the source. Implementation for ASM was overly complex, and an alternative was proposed. SSM is source specific multicast, where a subscriber joined a (group, source) tuple, to receive only from that source. SSM eliminated many of the issues associated with ASM for wide-scale deployment. It also signaled a new focus for multicast, which is one-to-many communication. ASM supported both one-to-many and many-to-many, while SSM is targeted particularly for the former. In other words, present day multicast is optimized for single source applications like webcasting. SSM was implemented via IGMPv3, replacing the ASM-based IGMPv2.

To expand beyond the MBONE, multicast needed a proper methodology for inter-domain routing. While BGP defined policy and reachability between domains for unicast traffic, it is not valid for multicast. In response, MBGP (Multicast Border Gateway Protocol) was developed.

Reliability was another concern, since multicast is implemented over UDP/IP. Needing a reliable equivalent of TCP, RMTP (Reliable Multicast Transport Protocol) was developed.

By 2001, these standards replace the older inferior protocols. By then, however, network operators associated multicast with problems experienced with the older protocols on the MBONE.

2.3 Multicast Issues

Given the obvious bandwidth savings, and a model seemingly naturally designed for content distribution, it would appear that multicast would be a heavily-used transport method in the currently media and content rich Internet. However, that is not exactly the case, for several reasons.

The biggest and most important of these is the limited penetration of multicast across the body of Internet's router infrastructure. As an optional feature, whose implementations were still maturing, it was frequently ignored in network deployments. ISPs became reluctant to support multicast, partially due to the added configuration and maintenance complexity (sometimes perceived to be greater than it was). The problem was exacerbated by the "dot-bomb" crash of 2000, as a result draining funding from a large number of leading-edge initiatives, significantly slowing down progress for new technologies, including multicast (thus the dissolution of partnerships like the IP Multicast Initiative).

The situation created a Catch-22. Customers did not sufficiently demand it because there were few providers. There were few providers because there were few potential customers. There were few potential customers because not enough ISPs supported multicast. In essence, it is an all-or-nothing situation. Like other technologies that require end-to-end adaptation for full functionality (IPv6 faces a similar difficult), multicast faces a significant barrier to relevance.

Therefore, multicast grew very slowly, first through educational and research network, then through early adopter ISPs. The rest of the Internet was cut off from its advantages, other than through tunneling or other workarounds. The current penetration is estimated to be somewhere between 5% and 15% [13].

3.1 Alternate Content Delivery Approaches

Meanwhile, content providers were faced with the problem of distributing content to large amounts of consumers. Naturally, the problem lay in the bandwidth – unicast sessions for popular content could quickly overwhelm a centralized source, particularly for data-intensive content such as audio and video, being accessed concurrently by many users.

One approach was the creation of CDNs – Content Distribution Networks. Essentially, CDNs work by replicating content among multiple servers located in distinct network regions, therefore splitting traffic along different routing paths, reducing load on any given link. This method was effective for static content, which could be distributed any number of ways to the servers, and benefited from existing effective caching strategies. The problem lay in delivering real-time content, such as streaming video. Reliability and synchronization between the edge servers required additional traffic, reducing some of the bandwidth benefits, and introducing additional failure points in the path to the end host. Another consideration is the additional cost of extra hardware.

3.2 P2P Overview

The content distribution method gaining most attention (and bandwidth) at the moment is P2P, or peer-to-peer networking. P2P establishes networks made up of many small servers (peers) that share and replicate content. Often, these servers are end hosts (customers) themselves, thus bandwidth and operating costs are pushed down to the users.

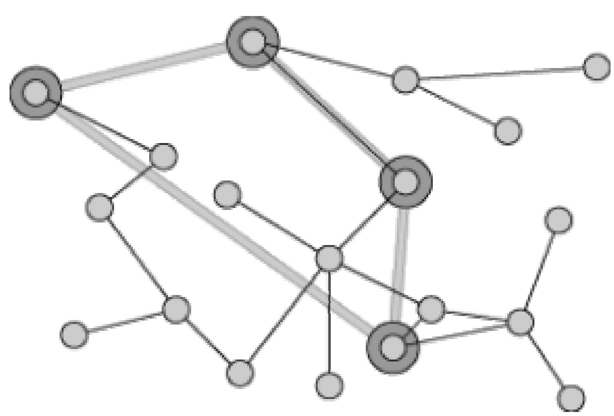

Figure 3: Nodes forming an overlay P2P network.

Typical peer-to-peer networks, including most of those that made the technology popular, such as Napster and Gnutella, are made up of many nodes sharing the same hierarchical level (or lack of thereof). The nodes are organized in an overlay network, in order for the member nodes to route information to each other. Frequently, the P2P establishes multicast over the overlay network – creating application level multicast. Inter-node routing is done via DHTs (Distributed Hash Tables), CANs (Content Addressable Networks), or similar approaches.

3.3 P2P Advantages

The killer app for P2P was of course file sharing. Starting with Napster as the first popular P2P network, millions of users have used the technology to share files, the vast majority of which has consisted of copyrighted music and later film content.

P2P was ideal for this because of its distributed nature. Huge file servers with expensive network links were unnecessary – the users themselves provided storage and connectivity. Large concurrent downloads became possible because content was replicated across multiple nodes. Replicated repositories on the Internet, commonly called mirrors (basic CDNs) became an obvious target for P2P distribution – an example is the BitTorrent P2P, which hosts many Linux distributions, a large and popular download that would cost a centralized distributor a lot of bandwidth.

Resiliency and fault tolerance are additional benefits of P2P. Because there is no "head," or root, in the hierarchy, it is more difficult to bring down a network in case of accidental failure or malicious attack. And scaling is easy, since peers can join and leave at will. In addition, the more peers join, the more the total available bandwidth, so supply naturally scales with demand.

P2P further provided other benefits of decentralization, such as certain legal protections – the reason illegal content was so popular. More recent P2P implementations such as BitTorrent and Freenet have gone to extensive lengths to separate and obfuscate the node from its content, to limit the responsibility of the node's owner.

3.4 P2P Shortcomings

While peer-to-peer structures have proved highly successful for file delivery, they do suffer from several disadvantages.

A major one is that, due to their popularity, P2P traffic is quickly becoming a significant share of all traffic traversing the Internet, and has thus caught attention of ISPs concerned about the load on their networks. With high-bandwidth edge (end-point) hosts, thanks to the spread of broadband, sharing content 24/7, and the rise of various legal issues surrounding P2P use, many ISPs have moved to limit P2P traffic. A whole new industry niche has spawned for companies like P-Cube that allow ISPs to identify and control P2P traffic; in the meantime, media companies have colluded to attack P2Ps indirectly by flooding them with corrupted content, further increasing the load on participating nodes.

Traffic use by P2P networks is high because, while at the level of the overlay network the routing is efficient, at the actual IP network layer, identical data can saturate the links for network peers. Indeed, research has been dedicated to make P2Ps more like multicast in order to reduce overall IP traffic, such as the MULTI+ protocol described in [1], or PeerCast in [2], or Nemo in [5]. But even approaches such as the popular ESM (end system multicast), developed in Carnegie Mellon, suffer performance penalties, due to the fact that "End System Multicast introduces duplicate packets on physical links, and incurs larger end-to-end delay than IP Multicast." [17] However, such complications would be unnecessary if IP multicast was utilized, substantially reducing the complexity of P2P communication as well as the corresponding bandwidth.

Because P2P technologies exist at the application level, there is a natural inter-communication boundary, forcing users to pick and choose implementations based on needs and features, rather than having a standards-based implementation.

Quality of service is a serious issue, since P2P relies on the kindness of strangers, so to speak. This problem is exemplified in the failure of vTrails and their Full Duplex Packet Cascading technology. FDPC attempted to approach the efficiency of IP multicast for streaming data by arranging peers in a hierarchy according to their bandwidth, and relying on them to retransmit data to the lower tiers. This is a problematic approach for a broadcaster, since it heavily relied on node over which it had little or no control, thus making it impractical to offer any QoS guarantees to the customer. Even with broadband, many of the end users are allocated asymmetric access by their ISPs (for example cable and DSL), with significantly lower upstream bandwidth – the unequal capabilities hurting the overall efficiency in the P2P network. This was shown in [6], showing an approximate 4.5/1 downstream/upstream ratio, and determining that ESM design must take this into account or suffer efficiency consequences.

The above hints at one type of communication that P2P cannot handle as well as multicast – one-to-many (such as web broadcasts). Multicast is optimized to place low stress on the content provider while giving it control over the distribution; P2P is optimized for peers to share content In other words, the "socialist" structure of peer-to-peer is non-optimal for the "totalitarian" structure demanded by content broadcasters.

4.1 So Why Multicast?

Even given the difficulties and slow adaptation by network operators, it is unlikely that multicast will be abandoned.

As it is, multicast has a niche in corporate intranets. Since the corporate network is typically high-speed and its operators have end-to-end control over the routing infrastructure, multicast is easier to enable. On the hardware level, it is not difficult – Cisco for example offers multicast across its family of network hardware, and on the software side, multicast is supported by a range of large providers such as Novell, Microsoft, Sun, etc. Thus, applications that can naturally benefit from multicast – teleconferencing, collaborative software, etc, have seen multicast implementations in the intranets. Examples are Microsoft's NetShow (and its ConferenceXP initiative) or Marratech's E-meeting Portal, as well as smaller applications like Network Assistant.

As we mentioned, the great barrier to multicast's global penetration is the chicken-and-egg problem bridging providers and consumers. Usage patterns today however are providing hospitable conditions for intensified interest in multicast This is a two-fold situation:

- increased availability of high-bandwidth, streaming content: audio (internet radio), video (news, entertainment), etc.
- increased consumer bandwidth due to the rise in broadband access (this summer, broadband use in the US overtook dialup for the first time)

This combination is rapidly driving up bandwidth requirements for content providers. Whereas a couple of years ago few people would utilize the higher-bandwidth content, now most are able (and more than willing) to do so. High-rate streams (up to 300KBps and higher) are becoming common, and during peak times, their unicast consumption can overwhelm the provider, such as in May of 2000, when two million Internet users tuned into the Victoria's Secret fashion show broadcast – even if only the 56K stream was utilized (for the lucky multicast users, a vastly superior 700K stream was available), that would add up to 300 Gb/s, a huge expense for the content provider, especially for a 30 minute event (if all users were tuned in for the whole show, that is 540 Tb of bandwidth) [7].

Another example illustrating this situation is BBC.co.uk, BBC's content provider on the web. In their own words,

> The BBC has been streaming since 94/95 and has coped with growing volumes of unicast traffic that are now limiting our use of content.
>
> So far we've catered for narrowband users with peaks of ~45K streams generating 2.2Gbit/s of traffic.
>
> We've avoided content that would greatly exceed that to avoid user disappointment though it's not been entirely problem free below there. [10]

Thus, this summer during the 2004 Olympics coverage, in order to alleviate the expected load (of up to 10Gbit/s), BBC offered multicast streams (of superior quality, at 370Kb/s, vs. 225Kb/s for unicast streams) to those UK customers whose ISPs peered with BBC. This was possible thanks to the rise in broadband use, and thus a rise in ISPs willing to route multicast content. As a result, BBC will be seeking to further exploring multicast streaming, and is actively seeking peering partners to extend its reach beyond the UK.

There are other recent examples of an escalation in the use of multicast for high-quality content delivery. In September, Path 1 Network Technology partnered with Modulus Video to produce broadcast-quality streams encoded with the recent MPEG-4 AVC format over IP networks using IP multicast. Multicast Technologies has already been providing DVD-quality multicast channels since 2002. World Multicast China announced a deal with Microsoft to bring its customers the first direct-to-home multicast IP television. In Japan, in July of this year, Plala Networks and Online TV launched 4th MEDIA service, a multi-channel broadcasting and video-on-demand service with over a 1000 titles across 40 channels, using the world's first commercial implementation of IPv6 multicast [8]. As the World Multicast CTO remarked, "IP Multicast is the only technology built-into the current inter-domain routing infrastructure that addresses mass media on the Internet."

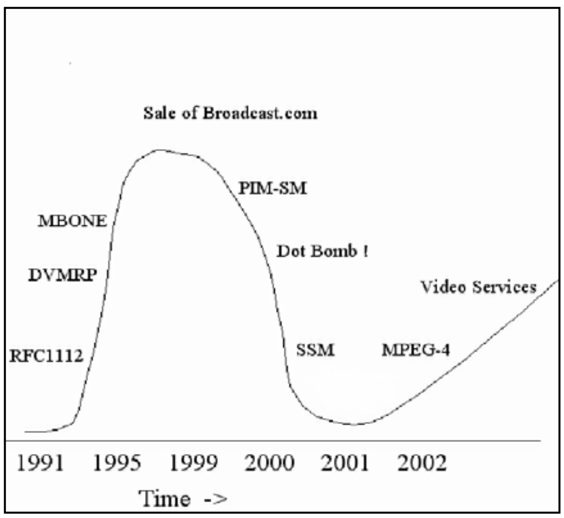

Figure 4: A potential path for multicast [13].

In essence, multicast provides the capability for broadcasters to deliver, for the first time, audio/video content over the web that could compete with "natural" multicast delivery like cable or satellite TV. Unicast technologies simply cannot scale, as the BBC case shows, to adequately deliver such content to a large consumer population. Nor is the application limited to the fixed Internet – as [11] shows, IP multicast is a viable mechanism for streaming over wireless networks. And [12] outlines the benefits of an IP multicast-enabled network for streaming content over 3G wireless providers. Indeed, IP multicast is part of the CDPD (Cellullar Digital Packet Data) standard already deployed by wireless providers such as Alltel.

As a result, high quality streaming content could be the "killer app" that accelerates multicast development and propagation. It gives both providers and consumers reason to demand multicast capability from the network intermediaries, thus giving the network intermediaries reason to enable multicast on their networks, with the added benefit of financial savings due to lower network utilization. It's not a coincidence that companies like Cisco and SkyStream offer IP multicast-based video servers targeted at broadcasters. The next logical step would be partnerships between ISPs and media companies to provide content to the ISP users – with multicast enabling the ISPs to do so.

So perhaps the slowdown in multicast interest in the early 2000s did not signify the end for multicast, but merely a necessary step in its evolution. Figure 4 illustrates a potential take on that path.

4.2 P2P and Multicast

As mentioned before, a lot of P2P technology relies on establishing application-level multicast, such as ESM. Much of P2P research relies on multicast behavior – for example, the application of P2P towards massively multiplayer gaming (MMG), as described in [4].

And while several attempts at commercial P2P approaches for streaming content have been made (several, like vTrails, failed), the approach presents a much more complicated and expensive solution to a content provider than would a multicast-enabled infrastructure. Figure 5 shows the structure of ChainCast, a P2P streaming provider – notice the number of different elements, each of which would require additional investment.

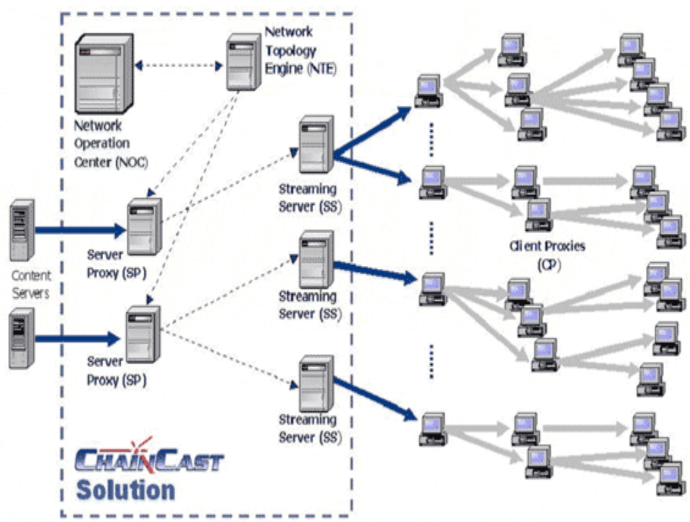

Figure 5: Infrastructure for ChainCast, a P2P streaming solution. [15]

The natural question to ask is whether P2P and
multicast are truly competing technologies, or complimentary – if multicast is generally available at the routing layer, P2P can benefit by building directly on top of it, simplifying its architecture. After all, without multicast, P2P still suffers from higher bandwidth utilization and higher complexity.

This is true for CDNs as well – native multicast allows CDNs to be more efficient, as was shown for example by the MC-A-RONI project by Eurescom [9]. In general, native multicast allows multi-target communication, whether by CDNs or P2P networks or else, to have simpler implementation (letting the corresponding protocols govern communication) and save bandwidth.

4.3 The Future of Multicast

As the partial penetration and limited applications described above have shown, multicast is not yet an ubiquitous, indispensable content delivery method. But indications are that its role will continue to expand, and with the advent of broadband access and the availability and demand for high-quality audio and video content, it may yet become a dominant for of communication.

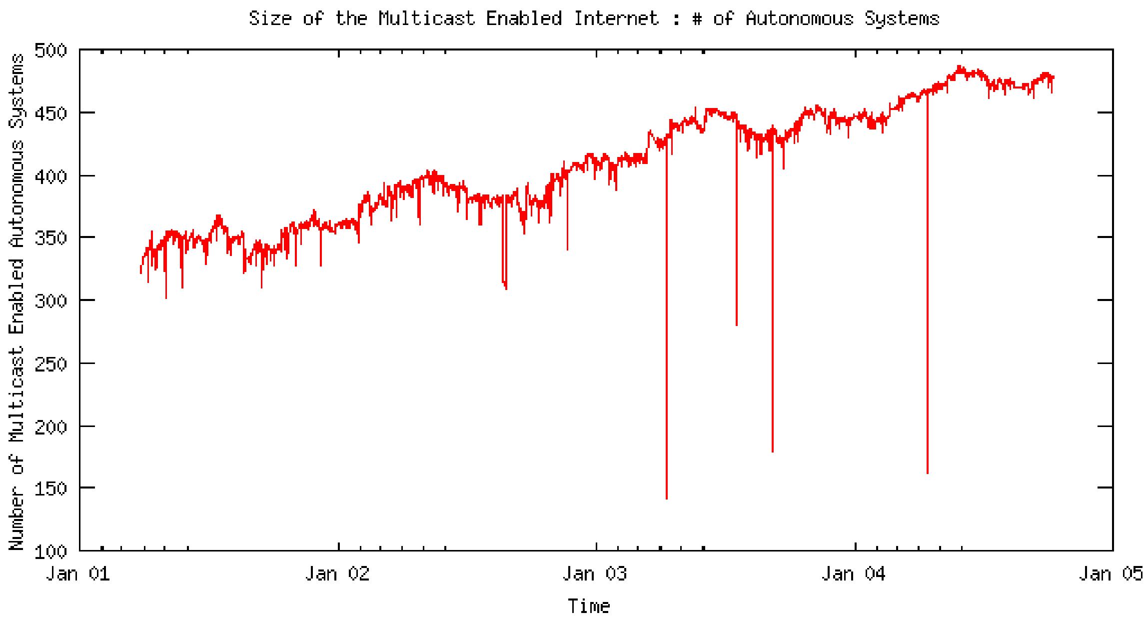

Figure 6: The number of multicast-enabled Autonomous System, as seen from AS 16517 – Multicast Technologies [18]

Figure 6 plots the number of multicast-enabled Autonomous Systems over the last four years, indicating a steady rise. While the Internet itself has been growing as well, multicast has not remained stagnant. Multicast is enabled for several large ISPs such as Sprint and WorldCom

Multicast is already supported across the spectrum of widely used media players, including those from Apple, Microsoft, and Real Media.

Furthermore, all indications are that broadband, and the speed of broadband, will continue to rise. For example, presidential candidate John Kerry included in his plan a national broadband policy seeking to increase broadband coverage as well as the average bandwidth [14]. All of this will certainly increase the load on content providers and network operators, further increasing the potential benefits of multicast deployment (already made simpler via standard network products from companies like Cisco).

What MBONE used to be to multicast research, Internet2 is now. Internet2 is a large consortium of universities, industry, and government working groups, with its own proving ground for high bandwidth technologies – the Abilene network, a high bandwidth backbone (which has multicast peering with several other large networks). Multicast is one of the major research areas of Internet2, and many of the major industry players (IBM, Cisco, Microsoft, etc) utilize Internet2 for multicast-enabled products.

Multicast is also a part of IPv6 development, with work on the relevant protocols, such MLD (Multicast Listener Discover Protocol) – the successor of IGMP. The focus in multicast research is on SSM – not only is SSM dependent on a smaller set of protocols than ASM, thus easier to implement, but as we've shown above, the industry interest in multicast is heavily in the one-to-many arena, thus favoring single source multicast.

5. Conclusion

In evaluating the evolution, functionality and competitive advantages of multicast and P2P, we outlined the positions these technologies hold on the spectrum of content delivery. We looked at the reasons for multicast's slow growth, and P2P's rapid rise, as well as the differences in their core competencies. We concluded that multicast has sufficient promise and proven advantage for certain types of content delivery, and recent escalation in commercial rollouts implementing multicast signify that its prime has not yet been reached. It is entirely possible that with continuing escalation in availability and demand of high-quality audio and video content, multicast will reach the saturation threshold in network penetration, at which point it can reach its potential in becoming a standard in content delivery.

6. References

1. Luis Garc´es-Erice, Ernst W. Biersack and Pascal A. Felber, *MULTI+: Building Topology-Aware Overlay Multicast Trees*, Institut Eurecom, Sophia Antipolis, France, April 2004.
2. Jianjun Zhang, Ling Liu, Calton Pu, Mostafa Ammar, *Reliable Peer-to-Peer End System Multicasting through Replication*, Fourth International Conference on Peer-to-Peer Computing (P2P'04), August 2004.
3. Thomas Karagiannis, Andre Broido, Nevil Brownlee, kc claffy, Michalis Faloutsos, *Is P2P dying or just hiding?*, IEEE Globecom 2004 - Global Internet and Next Generation Networks, Dallas, Texas, November, 2004.
4. Bjorn Knutsson, Honghui Lu, Wei Xu and Bryan Hopkins, *Peer-to-Peer Support for Massively Multiplayer Games*, INFOCOM 2004, March 2004.
5. Stefan Birrer, Fabián E. Bustamante, *Resilient Peer-to-Peer Multicast from the Ground Up*, Network Computing and Applications, Third IEEE International Symposium on (NCA'04), September 2004.
6. Karthik Lakshminarayanan, Venkata N. Padmanabhan, *Some Findings on the Network Performance of Broadband Hosts*, ACM/USENIX Internet Measurement Conference, Miami, October 2003.
7. SprintLink FAQ, *http://www.sprint.net/multicast/faq.html*
8. IPv6Style, July 12, 2004, http://www.ipv6style.jp/en/news/2004/0712_plala.shtml
9. MC-A-RONI - MultiCast Deployment - A Practical Handbook for Networks and ISPs, http://www.eurescom.de/public/projects/P1100-series/P1114/
10. BBC, http://www.bbc.co.uk
11. M. Hauge and Ø. Kure, *Multicast in 3G Networks: Employment of Existing IP Multicast Protocols in UMTS,* WoWMoM'02, Atlanta, USA, 2002.
12. *The Future of All-IP Broadband Wireless Mobile Networks*, ASWN 2004 Presentation, August 2004.
13. Mark Sutton, *The Case for Multicast*, Big Geek Program 04, January 2004.
14. John Kerry's Plan to Create Millions of High Wage Jobs in Industries of the Future, http://www.johnkerry.com
15. Robert C. Chalmers and Kevin C. Almeroth, *A Characterization of Multicast Efficiency*, http://www.nmsl.cs.ucsb.edu/mwalk/index.html
16. ChainCast, http://www.chaincast.com
17. Yang-hua Chu, Sanjay Rao, Hui Zhang and Srinivasan Seshan, *End System Multicast*, http://www.cs.cmu.edu/~sanjay/EndSystem/endsystem.html
18. Multicast Technologies, http://www.multicasttech.com