Multicast as a Viable Commercial Enterprise

Abstract
In our lifetimes, we have seen constant growth in the computer industries. Research and development of new protocols and novel ways of computer interconnections have been successful in permeating into everyday use. By the early 1980's, computer networks employing unicast communication protocols (one to one) were seen deployed across all types of networks. At the same time however, multicast was seen only in limited release in local area networks such as Ethernet and Token Rings. Twenty years later the scenario for multicast does not look much better. Even with commercial routers shipping with multicast support, only very few commercial applications of multicast are in use. In this paper we intend to show why the uptake on multitask is so painfully slow, and why, in its current incarnation, multicast is not ready of worldwide commercial deployment.

1. Introduction

The motivation behind multicast (one to many) is clear. In today's world of ever-changing content and end-user's broadband connections, traditional unicast (one to one) models of content delivery are terribly inefficient. For example, for a provider to handle all requests for a video file of considerable size using unicast, it must transmit a separate copy to each user. Since bandwidth is a limited, metered resource, multicast attempts to remove the need for this inefficient use by allowing shared delivery down shared spanning trees. It would appear then, by all accounts that multicast, in theory, seems like a wonderful technology that has many potential benefits to the internet services world. However, twenty years after its introduction, there are still not compelling reasons for the widespread deployment of multicast. With the fast, ever-changing pace of technology, there have to be reasons why a perfectly good theoretical idea has not seen its full potential.

In the very beginning of multicast's inception, multicast was in limited release – seen only in a few local area networks. However, those networks connected with bridges and routers did not support multicast data delivery. It would not be until the late 1980's with the introduction of multicast extensions to unicast routing mechanisms that multicast became a viable alternative over datagram-based inter-networks. The first "wide-spread" test-bed deployment of multitask was born – the MBone or Multicast Backbone. In 1992, the MBone had its event, streaming audio/video to about 20 sites worldwide. [5] The growth of the MBone provided impetus for router manufacturers to include multicast support in their routers and also shed some light into the then current multicast protocols. The following sections give a brief overview of the many multicast algorithms in use today, as well as examine the reasons why, even with router support, the default settings have multicast disabled.

We begin in Section 2 with a brief overview of multicast, describing its properties and features, and how these have potentially beneficial impacts on the internet-using world. Section 3 describes sizeable obstacles that the deployment of multicast has met, and why, in its current incarnation, infrastructure players will have a difficult time seeing the deployment of a viable multicast solution. Finally, we attempt to identify and dispute counter arguments in Section 4.

2. The Multicast Technology

There is an overhead associated with providing an acceptable level of service using multicast. Routers must support some multicast routing protocol, and clients must also implement support for joining/leaving and sending to/receiving from multicast groups. IP Multicast has a variable syntax and very loose semantics, which we will term the Multicast Model, and the Multicast Architecture.

2.1 The Multicast Model

IP Multicast is based on an "open service" model. In short, this means that it implements a very loose interpretation of access control. While "membership groups" exist, membership control is non-existent. Any client wanting to join a group can, simply by sending an IGMP (Internet Group Management Protocol) join message. Coupled with one of many of the routing protocols, the number of members subscribed at any one point in time is uncertain. Accordingly, the status of any sender or receiver is unknown due to the dynamic nature of joins and leaves. Any member or non-member with the desire to can send to any multicast group. We will see later why these semantics, along with some other properties of multicast are quite desirable from a provider's point of view.

In keeping with the "open-ness" of multicast, address allocation protocols are not specified. Any sender can take any multicast address, leading to multiple senders with the same multicast address. Adding to the uncertainty of the model is that based on the multicast routing protocol used, these multiple senders may or may not be on the same shared spanning tree – a fact that may cause undesirable side effects for any given sender/receiver pair.

What lies at the heart of the multicast model are the various routing protocols. The main difference among the algorithms is the type of spanning tree that each produces. DVMRP, MOSPF and PIM Dense Mode create trees by finding shortest paths between sender and receivers. PIM Sparse mode, CBT, OCBT and HIP base their trees on shortest distance between senders and a "central core" that all senders share. The potential for different routers to employ different trees leads to inconsistent views of multicast paths from sender to receivers. [3]

2.2 The Multicast Architecture

The most prevalent multicast combination in existence is comprised of IGMP v.2, DVMRP, MOSPF and PIM Sparse mode. These protocols alone, however, are only capable of multicast routing within an Autonomous System. Beyond the confines of these administrative domains, a protocol such as Multicast Source Discovery protocol or Multicast Border Gateway Protocol are necessary, with the former being utilized more frequently.

The sequence of events followed in multicast communication occurs generally as follows [2,3]:

1. IGMP messages are used by receivers to announce their interest in joining a particular multicast group. (IGMP v. 1 was used with DVMRP, the first routing algorithm developed, while IGMP v. 2 is the current IETF standard)
2. Edge routers receive this request and use the multicast routing protocols to shape a multicast spanning tree through the Internet.
3. When a host wants to leave a multicast group, it need not do anything
4. The multicast router sends queries out on each interface periodically to determine if interested hosts still exist for a given multicast address. If it does not receive a reply, it assumes no more members exist and prunes the appropriate portion of the tree.
5. On receipt of a query, a host can send a reply to tell the router it is still interested in the any given multicast group.

Protocols like DVMRP (including PIM dense mode and MOSPF) are flood-and-prune protocols. They flood routers with multicast packets and remove sections of trees based on a set of conditions. Inherent in any flooding algorithm is limited scalability. Flood-and-prune algorithms are not feasible to deploy across autonomous systems. They would bring the Internet as a whole down to a grinding halt by making routers respond and process multicast datagrams. The sparse mode protocols based on central cores have been introduced to alleviate the scalability problems in the flood and prune algorithms. However, these new algorithms introduced a new set of issues that are still not adequately solved.

3. Deployment Issues

As previously mentioned, multicast is at its core a very good idea. It allows for bandwidth savings, and groups senders and users into logical groups for the transmittal of data. Widespread adoption has been very slow and many service providers have found alternative ways to serve information without multicast protocols. Multicast is generally an all-or-nothing proposition – it can only work if all of the infrastructure players agree to deploy multicast solutions. After all, what good is multicast to commercial entities if support ends at its own domains' boundaries? What follows is an overview of the many issues that need to be solved before multicast can become a viable solution for all those involved.

3.1 Multicast Uses

While generally, a good idea, it is unclear whether multicast has strong commercial value. In a capitalistic society, commercial requirements need to be well defined, and a robust implementation structure implemented (or if not implemented, at least well designed). From the standpoint of the end user receiving data, he could care less whether his data comes as a multicast packet or unicast packet. Therefore the

impetus that drives multicast deployment must come from the service providers – wanting to save bandwidth.

There are only a few obvious uses for multicast are either one to many or few to many applications. These include [3]:

1. Audio/Video distribution – a very bandwidth intensive application could benefit from not having to send unicast packets to all subscribers.
2. Push applications – allow for subscriptions to information channels that push information at regular intervals (Even if a user is not present/reading)
3. Audio/Video conferencing – the defining moment in the MBone test bed. Multicast would allow a significant reduction of bandwidth requirements of users on the same shared spanning tree.
4. File transfer – similar to A/V distribution, multicast can help in the distribution of files from one/few to many.

We shall revisit these applications a little later to see if and how multicast could be a viable solution.

3.2 Multicast Requirements

There is a set of requirements that multicast protocols must meet in order for them to be attractive to their implementers. Unfortunately the current model and architecture of multicast do not provide easy solutions to these. And while much research has gone into finding ways to fulfill these requirements, current multicast implementations have a long way to go.

Since multicast really has no benefit to the end user, all the options and features of unicast should have an acceptable counterpart in multicast. In other words, services deployed using multicast and those deployed using unicast should be transparent to the user. End-users need to have access to all multicast services, meaning ISPs must provide a way for inter-domain multicasting. Conversely, multicast infrastructure components should be easy to install, and should impart as little inconvenience to ISPs and their customers as possible. In order for ISPs to have sufficient motivation to deploy multicast, there should be some way in which they can monitor usage statistics and other network management tasks on multicast clients. More importantly, however, is the ability for multicast to match financial and revenue generating models. The aforementioned loose access control may deter many companies from accepting multicast as an acceptable solution. Without the ability to manage, enumerate and restrict both senders and receivers, a multicast implementation could very quickly turn into a free-for-all anarchy – a situation no entity is likely going to want to accept responsibility for. Also, before multicast can be even considered viable, some guarantee to providers must be made that their multicast address will be unique, at least for the duration of a particular session. Finally, multicast users should be able to expect the same reliability as unicast users.

3.3 Multicast Deployment

It should be duly noted, that even if the above-mentioned requirements are met, there exist many more obstacles in the full deployment of multicast. The multicast architecture arguably requires more setup and administration than its unicast counterpart. But even after implementing multicast, some carriers have reported that the current multicast architecture is unstable.

There is a current pattern by which most ISPs adhere when deploying routers that maximizes the value of each router. They begin by deploying the newest (and most expensive) routers near the backbone, and as new technology becomes available, they begin moving the routers closer to their customers. This provides a high level of service near the backbone and allows full use of routers as they age. Since many old routers do not support multicast, implementing multicast would require that ISPs overhaul their entire infrastructure (rather than piece meal). They would not be able to write off router depreciation for tax purpose and would be effectively discarding a large number of routers. There is also the idea that most routers deployed near the backbone are less intelligent and thus can handle large amounts of traffic. The overhead associated with processing multicast datagrams would therefore diminish these routers capacity.

Also of great concern to ISPs are inter-domain dependencies. Much of the multicast infrastructure requires trust among ISPs. It is entirely possible for the central core of a spanning tree to exist outside of one particular

ISP's domain. This being true, the ISP has little control over the administration of that tree. ISPs also have no incentive to handle traffic for multicast groups for which they have no senders or receivers. All of these factors will impede any progress that is made to fulfill the requirements in the previous section.

3.4 Security Considerations

A brief look into the security of multicast already reveals much inadequacy. Since multicast trees can exist in different autonomous systems than both sender and receiver, nothing guarantees the security of the data or the security of the routes. The loose semantics of access control allow anyone, with the desire to, to join a group and receive any data packets destined for the multicast group. And since no control can be exerted to who can send, any malicious user could flood a multicast group with packets, effectively causing a Denial of Service attack. Also inherent in the multicast architecture is the possible exposure of network topologies to outside parties. IGMP packets and responses to IGMP queries may reveal information otherwise kept private. Lastly, due to the absence of address allocation, data from different multicast sessions may become interleaved and render all received data unusable. Currently, there are no provisions within the model or architecture to allow multicast traffic without exposing the network to the above vulnerabilities. [1,3]

4. Counter Arguments

Even with the evidence presented before them, there exist zealots for whom the prospect of worldwide deployment of multicast is possible today, with the current model and architectural designs. For them, the absence of global addressing techniques, the presence of security holes, the difficulty of network and user management, and no observable benefit for the end user are overlooked and replaced by several simple, but flawed ideas.

4.1 Multicast has been deployed and shown to be successful

There are implementations of multicast that work, but their numbers are very slim – for reasons already mentioned in this paper. The Multicast Backbone has been a great asset for many who research the possibilities of multicast. However, it is unclear as to the how multicast would fare in a global environment. Many intranets have deployed multicast for private use, and we have seen already that multicast is a fine solution for internal networks (small autonomous systems). The Seismic Research Symposium has deployed a reliable multicast network [1] that includes satellite links as well as traditional LANs interconnected. However even this implementation is only on a private network where system/network administrators have full control over multicast addresses and host software. None of the problems encountered in global multicast deployment are ever encountered.

4.2 Reliable multicast is available and can mimic unicast transparently

Reliable multicast is available. The problem is *many* versions of reliable multicast exist. [1] Perhaps in preserving the end-to-end semantics of multicast, designers left out the provision for reliable communications. The result is application specific measures to ensure reliability. Since many people's ideas of reliability may differ, no true standard currently exists that is analogous to TCP on unicast connections. TCP is deployed as a standard protocol across most operating systems, and is well defined. Until such a protocol is developed for multicast platforms, true transparency cannot be attained by multicast applications.

4.3 Multicast is cost effective

Perhaps sometime in the distant future, long after the global implementation of multicast, multicast will become cost effective. At the moment however, the cost of implementation, maintenance and overall management far exceeds any bandwidth savings. The total cost of a network service like multicast is the sum of all network related costs. (This includes router state, processing, signaling, and inter-domain routing scalability) [3] Ideally, in order for multicast to become cost effective, the cost of deployment must be offset by bandwidth savings. Normally, for unicast, as the number of clients increases, the cost to support the clients increases rather linearly. One would think that because of multicast's properties, that increasing clients in a multicast environment would show a decrease in cost/client. While this may be the ideal situation, currently, the addition of clients in the multicast environment can show a logarithmic increase, especially if the clients are in different autonomous systems. This would indicate that while multicast *could* save some bandwidth with a large number of clients, it *will* be quite a bit more expensive for fewer clients. Accordingly, ISPs would have little incentive to implement

multicast if multiple unicast sessions would do the job just as well. Multicast remains a very expensive proposition. [3]

Also uncertain with respect to the true cost of multicast is the idea of cost sharing. Since multicast requires that many autonomous systems cooperate in order to provide the specified network service, there is no clear way of dividing transmission costs of multicast packets. A study done on cost analysis of multicast systems [4] found that there exists no mechanism that is both efficient and budget balanced among all involved parties. Where the latter means that the sum of all charges over all parties covers the total incurred cost of the multicast service. If these two goals are at odds, there is little hope that infrastructure providers will opt for efficiency.

5. Conclusion

When one looks around and examines the many network services available today, multicast would seem to be like a relatively new and emerging technology – not a twenty year old one struggling to gain a foothold in the Internet world. There seem to be so many unanswered questions and unresolved issues that make multicast unfit for deployment on a grand scale. Even the valid reason of saving bandwidth and ultimately cost does not fare well with the current model and architecture. It would seem, for the moment and until whenever these issues are addressed, that multicast will be relegated to small pockets of deployment within single autonomous systems or administrative domains.

References

[1] Agarwal, D. and Lawrence, O. "Using Multicast in the Global Communications Infrastructure for Group Communications". 21st Seismic Research Symposium.

[2] Agilent Technologies, "Testing IP Multicast".

[3] Diot, C. et al. "Deployment Issues for the IP Multicast Service and Architecture". 2000.

[4] Feigenbaum, J. et al. "Hardness Results for Multicast Cost Sharing".

[5] Shi, S. "A Brief History of Internet Multicast".
http://www.arl.wustl.edu/~sherlia/thesis/chap1/node3.html