Tsunemasa Hayashi, NTT
Haixiang He, Nortel Networks
Hiroaki Satou, NTT
Hiroshi Ohta, NTT
Susheela Vaidya, Cisco Systems

Internet Draft
Expires: April 15, 2006

October 12, 2005

Issues Related to Receiver Access Control in the Current Multicast Protocols

<draft-ietf-mboned-rac-issues-01.txt>

Status of this Memo

Copyright Notice

Abstract

This I-D evaluates the extent to which current multicasting protocols can be used to address common requirements for commercial, large-scale IP multicasting. Four existing possible multicasting architectures (with or without some form of access or content control) are presented. Then each architecture is analyzed with respect to how it can or cannot satisfactorily address each issue. This I-D concludes that for many of these issues the possible architectures based on present standards as they now exist require non-standardized solutions to meet common use requirements. This I-D recommends for requirements to be defined that would set the groundwork for creating standardized ways to overcome these limitations.

Copyright Notice...1
1. Introduction..3
2. Definitions and Abbreviations.....................................4
2.1 Definitions..4
2.2 Abbreviations..4
3. Common use models and network architecture implications.........5
4. Issues in multicasting related to commercial and large-scale implementations..6
4.1 Access limits and resource issues...............................6
4.2 Capability to distinguish between receivers (end hosts)........6
4.3 Capability to distinguish between users (as opposed to merely hosts)...7
4.4 Channel "leave latency"...7
4.5 Surveillance of receiver by sender..............................7
4.5.1 Precise access log..7
4.5.2 How to share user information.................................7
4.5.3 Trustworthy logs to monitor user activity....................8
4.6 Notification to users of the result of the join request........8
4.7 Triple Play..8
4.8 DRM Protection...8
5. Description of existing architectures...........................8
5.1 IGMP/MLD...8
5.2 IGMP/MLD plus L2/L3 Authentication with Access Control Policy.10
5.3 Unicast Control with IGMP/MLD..................................11
5.4 IGMP/MLD with Multicast Encryption.............................11
6. Evaluation of architectures by issue...........................12
6.1 Access limit capabilities, compared by architecture...........12

6.2 Capability to distinguish between receivers, compared by architecture...13
6.3 Capability to distinguish between users, compared by architecture 13
6.4 Maintain guaranteed quality-level of data delivery (Voice, Video), compared by architecture...14
6.5 Fast leave for fast surfing capability, compared by architecture 14
6.6 Surveillance of receiver by sender, compared by architecture..15
6.7.Notification to users of the result of the join request compared by architecture...15
6.8 Comparison summary...16
7. IANA considerations...16
8. Security considerations.......................................16
9. Conclusion..16
Normative References...17
Full Copyright Statement...19
Intellectual Property..19
Acknowledgement..19

1. Introduction

The intention of this I-D is to initiate a discussion on the state of current multicasting protocols deployed for commercial, large-scale multicasting and their capabilities to provide receiver access control.

Existing IP multicasting protocols (as presented in Section 5) were designed to meet certain sets of requirements that do not necessarily include architectural considerations intended to support commercial services. This I-D presents a number of issues network providers may face when they attempt to apply current multicasting standards to commercial services. The extent to which existing multicast protocols can or cannot satisfactorily deal with these issues is explored. A few network models based on a range of different business models are presented as a basis for defining requirements.

Multicasting can be useful to make the network more scalable when a large volume of information needs to be distributed to a large number of receivers. However, multicasting according to current standards (e.g., IGMPv3[1] and MLDv2[2]) has drawbacks compared to unicasting in terms of its commercial applicability because of the insufficiency of access control and protect network resources against malicious use or accidents. In order to be applicable to large-scale commercial networks, multicast networks need to have the same capabilities which are currently supported by unicast networks. Such issues which are

important to commercial, large-scale implementations of multicasting are listed. Next, a few possible existing architectures used for multicasting with access control based on current standards are presented. Specifically 1) IGMP/ MLD, 2) IGMP/MLD with L2 Authentication with ACL 3) Unicast Control with IGMP/MLD and 4) IGMP/MLD with Multicast Encryption will each be presented and described. Each architecture is discussed with respect to the presented list of issues.

2. Definitions and Abbreviations

2.1 Definitions

For the purposes of this I-D the following definitions apply:

Accounting: actions for grasping each user's behavior, when she/he starts/stops to receive a channel, which channel she/he receives, etc.

Authentication: action for identifying a user as a genuine one.

Authorization: action for giving permission to access the content or network to a user.

Receiver: an end-host or end-client which receives content. A receiver may be distinguishable by a network ID such as MAC address or IP address.

User: a human with a user account. A user may possibly use multiple reception devices. Multiple users may use the same reception device.

Note: The definition of a receiver (device) and a user (human) should not be confused.

2.2 Abbreviations

For the purposes of this draft the following abbreviations apply:

ACL: Access Control List

CDS: Content Delivery Services

CSP: Content Service Provider

DRM: Data Rights Management

KEI: Key Exchange Identifier

NSP: Network Service Provider

QoS: Quality of Service

3. Common use models and network architecture implications

Issues such as user identification, access-control, tracking and billing are common requirements for commercial content delivery services (CDS) systems (and are important in many non-commercial CDS systems as well.) These same requirements should be met for CDS systems that employ multicasting.

In some cases a single entity may design and be responsible for a system that covers the various common high-level requirements of a commercial multicasting system such as 1) content serving, 2) the infrastructure to multicast it, 3) network and content access control mechanisms. In many cases however the content provision and network provision roles are divided between separate entities. The I-D draft-ietf-mboned-maccnt-00.txt provides more detail of the multiple versus single entity CDS network model.

As such it should not be assumed that the entity responsible for the multicasting structure and the entity responsible for content serving are the same. Indeed because the infrastructure for multicasting is expensive and many content holders are not likely to be competent at building and maintaining complicated infrastructures necessary for multicasting, many content holders would prefer to purchase transport and management services from a network service provider and thus share the infrastructure costs with other content holders.

Similarly commercial network service providers do not generally specialize in providing content and are unlikely to build and maintain such a resource-intensive system without a certain level of demand from content holders.

The business model of a single NSP providing multicasting services to multiple CSP has certain implications:

     -Need for user tracking and billing capabilities
     -Need for network access control and/or content access control
satisfactory to the requirements of the CSP
     -Methods for sharing information between the NSP and CSP to make
the above two possible

When the NSP and CSP are the same single entity the general requirements are as follows.

     -Need for user tracking and user-billing capabilities

-Need for access control and/or content protection at level the entity deems appropriate

In the next section issues in multicasting related to commercial and large-scale implementations are presented. Some presented issues are not pertinent to cases where the NSP and CSP are the same entity.

4. Issues in multicasting related to commercial and large-scale implementations

This section lists issues related to receiver access control in current multicasting protocols which are especially important to commercial, large-scale multicasting. More details concerning the requirements related to these issues are provided in a separate I-D draft-ietf-mboned-maccnt-00.txt[3]. References to that document are provided as applicable below.

4.1 Access limits and resource issues

For commercial applications of multicasting, network and content providers generally wish to be able to control the number of groups a host can access at the same time. Also the network provider may wish to limit the number of users accessing a multicast stream because of bandwidth and processing issues between the receiver and the multicast server.

With best-effort services (e.g. mail transfer, web surfing) strict network resource allocation is not necessary, but for services with a guaranteed QoS level (e.g. IP television, teleconferencing, VoIP) it is necessary to allocate sufficient bandwidth and server resources to each service. In order to guarantee certain QoS levels, it is important for network providers to be able to protect their network resources from being wasted (either maliciously or accidentally).

More detail on this topic is provided in I-D draft-ietf-mboned-maccnt-00.txt, section "Issue of network resource protection."

4.2 Capability to distinguish between receivers (end hosts)

Currently the sender cannot distinguish which receivers (end hosts) are actually receiving its information with existing protocols (IGMP/MLD.) The sender must rely on the information from the multicasting routers. This can be complicated if the sender and routers are maintained by different entities. There is currently no standard way to share such information.

4.3 Capability to distinguish between users (as opposed to merely hosts)

Many content providers would like to have detailed information on which users (as opposed to merely hosts identified by physical addresses, etc.) are consuming their content and information on their usage behavior. More detail on this topic can be found in I-D draft-ietf-mboned-maccnt-00.txt, section "User identification."

4.4 Channel "leave latency"

Commercial implementations of IP multicasting are likely to have strict requirements in terms of user experience. Leave latency is the time between when a user sends a signal that he/she wishes to "leave" a group and when the network recognizes the "leave."
A separate I-D draft-ietf-mboned-maccnt-00.txt provides more detail on this topic in the section "Channel 'leave latency'"

4.5 Surveillance of receiver by sender

4.5.1 Precise access log

It is necessary to precisely log information such as who (host/user) is accessing what content at from what time (join action) until what time (leave action). The result of the access-control decision (e.g. results of authorization) would also be valuable information.

4.5.2 How to share user information

For commercial multicast applications where NSP and CSP are different entities, there are a number of issues regarding how to share user information between the NSP and CSP. For example, which entities should be able to access which information relating to user-based tracking? What is the user identifier that can be used between the entities to distinguish among users, and which entities should be able to recognize this identifier? Another important issue is how the edge router should be able to access and then maintain user information. The current situation of present architectures is that only the NSP can get information about user activity, because user activities are only observable from join/leave information logged on edge devices which are under control of the NSP.

4.5.3 Trustworthy logs to monitor user activity

An important issue for commercial multicasting applications is how the NSP can get trustworthy data on user activity which may be needed for billing and statistics purposes. A standard way of logging user activity and protecting the integrity of the logs does not exist. Often network providers do not want to keep logs on untrusted user terminals which can be tampered with.

4.6 Notification to users of the result of the join request

It is necessary to provide information to the user about the status of his/her join request(granted/denied/other).

4.7 Triple Play

Ideally the NSP should be able to use the same infrastructure (such as access control) to support commercial multicast services for the so-called "triple play" services: voice (VoIP), video, and broadband Internet access services.

4.8 DRM Protection

Digital Rights Management (DRM) is important but out of scope of this I-D.

5. Description of existing architectures

In this section, existing architectures used for multicasting based on current standards are defined. In section 6 these architectures will be compared by the issues presented in section 4.

5.1 IGMP/MLD

Internet Group Management Protocol(IGMP) or Multicast Listener Discovery (MLD) are protocols for layer 3 management of multicasting. In IP multicast a receiver sends a request to a first-hop multicast router to join a particular multicast group. The router is then responsible for forwarding the appropriate data from the sender to the receiver.

```
  +----------+    +----------+   +----------+         +----------+
  | Sender   |    | Router   |   | L2SW     |         | Receiver |
  |          |    |          |<---------------1,JOIN--|          |
  |          |    |          |   |          |         |          |
  |          |--------------------------------2,Data->|          |
  |          |    |          |   |          |         |          |
  |          |    |          |   |          |         |          |
  +----------+    +----------+   +----------+         +----------+
```

For the sake of simplicity, the above diagram only shows the sequence of requests for a single receiver. When multiple receivers are requesting the same channel stream the data would be copied at the multicasting router to serve the multiple streams.

5.2 IGMP/MLD plus L2/L3 Authentication with Access Control Policy

With a basic implementation of IGMP/MLD implementation, no authorization is performed on the receiver. It is possible to combine an IGMP/MLD implementation with Layer 2 or Layer 3 Authentication to provide an access-control mechanism at the first point of attachment to the network, for example, using 802.X.

For example, a receiver may request to an L2 authentication server for access to the network. The authentication controller then queries the policy server with the receiver's credentials (such as IP or MAC address), and if the receiver is determined to be an authorized user of the network ("success"), the router downloads the ACL from the policy sever. For example, users which are not on the ACL are rejected. Then the Layer 2 Switch is directed to open a port for the receiver to send a join request to the multicast router. The router is then responsible for forwarding the appropriate data from the sender to the receiver.

Note: ACL is one existing method to realize an access control policy. Other methods exist.

```
     +----------+
     | Policy   |
     | Server   |\
     |          | \
     +----------+  \ 4,ACL Download
        |    ^      \
        |    |       \
        V    |        V
     +----------+    +----------+   +----------+          +----------+
     | L2       |    | Router   |   | L2SW     |          | Receiver |
     |          |    |          |   |          |          |          |
     | Auth.    |<----------------------------  1,Request-|          |
     |          |    |          |   |          |          |          |
     |          |--------2,Success------------>X(3,Auth)  |          |
     +----------+    |          |   |          |          |          |
                     |          |   |          |          |          |
     +----------+    |          |   |          |          |          |
     |          |    |          |<---------------5,Join---|          |
     | Sender   |    |          |   |          |          |          |
     |          |------------>x-----------------6,Data-->|          |
     |          |    |        | |   |          |          |          |
     +----------+    +--------|-+   +----------+          +----------+
                              |
                              V
```

Key:
 Auth: Authentication

5.3 Unicast Control with IGMP/MLD

The receiver first sends a unicast request to the sender which resides in the CP's domain. This method is the same as that used in unicast video-on demand (VoD) systems. If authorization is successful the sender sends the multicast address information via unicast. With this multicast address the receiver does a IGMP\MLD join as in described in 5.1.

```
     +----------+    +----------+   +----------+          +----------+
     | Sender   |    | Router   |   | L2SW   |          | Receiver |
     |          |    |          |   |          |          |          |
     |          |<------------------------------1,Request-|          |
     |          |    |          |   |          |          |          |
     |          |-------------------------------2,Success>|          |
     |          |    |          |   |          |          |          |
     |          |    |          |<--------------3,Join----|          |
     |          |    |          |   |          |          |          |
     |          |------------>x-----------------4,Data--->|          |
     |          |    |        | |   |          |          |          |
     |          |    |        | |   |          |          |          |
     +----------+    +--------|-+   +----------+          +----------+
                              |
                              V
```

5.4 IGMP/MLD with Multicast Encryption

With a basic implementation of IGMP/MLD, no data protection is performed on data sent to the receiver. No credential check is performed on the receiver and any receiver can receive and use the data. The IGMP/MLD with Multicast Encryption model assumes that the sender is sending encrypted data and that for this data to be useful to the receiver it must first request and receive a key from a group controller and key server that is synchronized with the content encryption occurring on the sender's data.

```
  +----------+    +----------+   +----------+          +----------+
  | G.C. &   |    | Router   |   | L2SW     |          | Receiver |
  |          |    |          |   |          |          |          |
  | Key S.   |<------------------------------1,Request-|          |
  |          |    |          |   |          |          |          |
  |          |-------------------------------2,Key---->|          |
  +----------+    |          |   |          |          |          |
                  |          |   |          |          |          |
  +----------+    |          |   |          |          |          |
  |          |    |          |<---------------3,Join---|          |
  | Sender   |    |          |   |          |          |          |
  |          |------------>x------------------4,Data-->|          |
  |          |    |        | |   |          |          |          |
  +----------+    +--------|-+   +----------+          +----------+
                           |
                           V
```

Key:
G.C. & Key S.= Group Controller and Key Server

6. Evaluation of architectures by issue

In this section the various issues raised in section four are analyzed by each of the architectures introduced in section five.

6.1 Access limit capabilities, compared by architecture

Comparison of currently available architectures with respect to limiting the access of multicast groups

- IGMP/MLD: It is not possible to limit data reception.

- L2/L3 authentication with access control policy:
With an ACL it is possible to limit access of multicast groups. However it should be discussed as to how scalable this approach is because configuring an ACL could be a labor-intensive task.

- IGMP/MLD with Unicast control
It is possible for malicious users to reconfigure the receiver's terminal to ignore the Unicast control. In this case, this maliciously reconfigured terminal could send a join message even if it is rejected by the network. In such a case, the ineligible receiver would be able to receive the multicast. As such, this method may not be strong enough to exclude ineligible access.

-Multicast Encryption:
It is possible for receivers to receive IP packets, even if they do

not possess the keys to decrypt them. A receiver may also be able to store such received data until they discover a way to decrypt it. Another disadvantage of this method is that network resources are wasted if an ineligible receiver receives an encrypted content even if they do not have a valid key.

6.2 Capability to distinguish between receivers, compared by architecture

Comparison of currently possible protocol-based solutions.

-IGMP/MLD:
The sender has no direct line of contact with the receiver and therefore cannot distinguish on a receiver-basis.   (If the interface is fixed to the receiver then the join-log can be used, but this would mean portability is sacrificed. Moreover, this method is not applicable to a case where the CSP and NSP are different companies because CSP cannot access this join-log.)

-L2/L3 authentication with access control policy:
At the moment of L2/L3 authentication it is possible to recognize receivers, but if there are multiple content service providers (CSP) a single L2 Authorization Server cannot distinguish among the CSPs. Therefore it would be necessary to gather the join logs. (If the interface is fixed then the join-log can be used, but this would mean portability is sacrificed. Moreover, this method is not applicable to a case where the CSP and NSP are different companies because the CSP cannot access this join-log.)

-IGMP/MLD with Unicast control :
It is possible to distinguish among receivers using Unicast control.

-Multicast Encryption:
If the Content Service Provider maintains the Key Server it is possible to distinguish on the receiver-level. If the Network Service Provider maintains the key server it is necessary to devise a method for the NSP to notify the CSP.

6.3 Capability to distinguish between users, compared by architecture

Comparison of currently possible protocol-based solutions:

-IGMP/MLD:
Since there is no user-based information, it is not possible to distinguish on the user-level.

-L2/L3 authentication with access control policy:
At the moment of L2/L3 Authentication it is possible to distinguish

on the user-level.

However it is difficult to combine user and group logs: it would be necessary to match user IDs from L2-Auth logs and group IDs from the Join logs to match users and groups.

-IGMP/MLD with unicast control :
Distinguishing by user is possible using unicast control.

-Multicast Encryption:
If the Content Provider manages the Key Server it is possible to distinguish the user.
If the Network Service Provider manages the Key Server it is necessary to notify the Content Provider.

6.4 Maintain guaranteed quality-level of data delivery (Voice, Video), compared by architecture

Comparison of currently possible protocol-based solutions:

-IGMP/MLD:
It is not possible to reject a user attempting to access even if there are not sufficient resources.

-L2/L3 authentication with access control policy:
The AAA server does not know whether there are sufficient resources or not. This method still can provide a guaranteed QoS if every channel has the same bandwidth and sufficient bandwidth are allocated to each user beforehand. However, it is not possible to provide a guaranteed QoS by comparing the available bandwidth and the necessary bandwidth upon each user's request.

-IGMP/MLD with Unicast control :
When the CSP and NSP are separate entities it is not possible for the CSP to make a proper authorization decision because only the NSP grasps the network resource availability.

-Multicast Encryption:
It is not possible to reject a user attempting to access even if there are not sufficient resources because the user can receive data even without a valid key.

6.5 Fast leave for fast surfing capability, compared by architecture

Comparison of currently possible protocol-based solutions:

-IGMP/MLD:
It is possible to track on a per host level (based on host address)

therefore fast leave for fast surfing capability can be achieved.

-L2/L3 authentication with access control policy:
It is possible to track on a per host level (based on host address) therefore fast leave for fast surfing capability can be achieved.

-IGMP/MLD with Unicast control :
Even if a quick leave is possible, changing to a new channel using Unicast Control is slow (latency problem).

-Multicast Encryption:
Even if a quick leave is possible, delivery of the Key Exchange Identifier(KEI) is slow.

6.6 Surveillance of receiver by sender, compared by architecture

Comparison of currently possible protocol-based solutions:

-IGMP/MLD:
With this protocol it is possible to separately log join and leave actions, but it is difficult to match these join and leave actions because analyzing the logs requires heavy computation (related to the scalability with millions of users).

-L2/L3 authentication with access control policy:
In this solution, the leave action is not recorded unless some additional mechanism such as IGMP/MLD snooping is used. In some cases, users disconnect their terminals without sending leave messages. In this case, it is not possible to determine when each user's entry in the ACL should be deactivated.

-IGMP/MLD with Unicast control :
In this solution the leave action is not recorded.

-Multicast Encryption:
If logs are recorded for each renewal of keys, then it is possible to track activity on a per-user basis. However if logs are only recorded per content data download then such tracking is not possible.

It should be noted that authentication of the source of each join/leave message is important.

6.7.Notification to users of the result of the join request compared by architecture

Comparison of currently possible protocol-based solutions:

-IGMP/MLD:
After the join it is not possible to notify the user of the result of the join request.

-L2/L3 authentication with access control policy:
After the join it is not possible to notify the user of the result of the join request.

-IGMP/MLD with Unicast control :
After the join it is not possible to notify the user of the result of the join request.

-Multicast Encryption:
After the join it is not possible to notify the user of the result of the join request.

6.8 Comparison summary

In this section a variety of existing architectures used for multicasting based on current standards were compared and evaluated. None of these architectures can sufficiently meet all of the common requirements for accounting, authentication and authorization in commercial, large-scale IP multicasting. Therefore it is recommended that framework(s) for sufficiently addressing such requirements be explored.

7. IANA considerations

This I-D does not raise any IANA consideration issues.

8. Security considerations

This I-D does not raise any new security issues which are not already existing in original protocols. Enhancement of multicast access control capabilities may enhance security performance.

9. Conclusion

Issues such as user identification, access-control, tracking and billing are common requirements for many content delivery services (CDS) systems. When CDS systems employ multicasting with architectures based on currently existing multicasting standards, it is often necessary to deploy non-standardized solutions to meet these common requirements. It is recommended that requirements be defined to serve as a basis for creating standardized ways to address the various issues discussed in this I-D which are limiting the application of multicasting especially to commercial, large-scale CDS services. Such requirements should take into consideration a range of possible architectures based on multiple business or usage models.

Normative References

[1] B. Cain, et. al., "Internet Group Management Protocol, Version 3", RFC3376, October 2002.

[2] R. Vida, et. al., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC3810, June 2004.

[3] Hayashi, et. al., "Accounting, Authentication and Authorization Issues in Well Managed IP Multicasting Services", draft-ietf-mboned-maccnt-req-01.txt, October 2005 [Work in Progress].

Authors' Addresses

Tsunemasa Hayashi
NTT Network Innovation Laboratories
1-1 Hikari-no-oka, Yokosuka-shi, Kanagawa, 239-0847 Japan
Phone: +81 46 859 8790
Email: hayashi.tsunemasa@lab.ntt.co.jp

Haixiang He
Nortel Networks
600 Technology Park Drive
Billerica, MA 01801, USA
Phone: +1 978 288 7482
Email: haixiang@nortelnetworks.com

Hiroaki Satou
NTT Network Service Systems Laboratories
3-9-11 Midoricho, Musashino-shi, Tokyo, 180-8585 Japan
Phone : +81 422 59 4683
Email : satou.hiroaki@lab.ntt.co.jp

Hiroshi Ohta
NTT Network Service Systems Laboratories
3-9-11 Midoricho, Musashino-shi, Tokyo, 180-8585 Japan
Phone : +81 422 59 3617
Email: ohta.hiroshi@lab.ntt.co.jp

Susheela Vaidya
Cisco Systems, Inc.
170 W. Tasman Drive
San Jose, CA 95134
Phone: +1-408-525-1952
Email: svaidya@cisco.com

Comments

Comments are solicited and should be addressed to the mboned working group's mailing list at mboned@lists.uoregon.edu_and/or the authors.

Full Copyright Statement

Intellectual Property

Expiration

This Internet-Draft will expire on April 15, 2006.

Acknowledgement

Funding for the RFC Editor function is currently provided by the

Internet Society.