Hiroaki Satou, NTT
Hiroshi Ohta, NTT
Tsunemasa Hayashi, NTT
Haixiang He, Nortel Networks

Internet Draft
Expires: December 25, 2006

June 23, 2006

AAA Framework for Multicasting
<draft-ietf-mboned-multiaaa-framework-01.txt>

Status of this Memo

This Internet-Draft will expire on December 25, 2006.

Copyright Notice

Abstract

This memo provides a generalized framework for solution standards to meet the requirements presented in draft-ietf-mboned-maccnt-req-04.txt, "Requirements for Accounting, Authentication and Authorization in Well Managed IP Multicasting Services". In this framework a user sends a request for multicast data to a network service provider. The network service provider selects the appropriate content provider to send the user's request. The request is sent by the network service provider to the content provider transparently in a way so that the network service provider and content provider do not need to know the corresponding user id for the same user in the other provider's domain. The content provider then responds with an indication of "success" or "failure" to the network provider and in the case of "success", the network provider may delivery the requested data to the user. The network service may base its decision to deliver such data to the user based on its bandwidth management policy. The framework is designed to be flexible and extendible, so it will be possible to implement partially enabled versions as well as fully enabled versions of the model. Also an additional entity may provide transit of requests between network service providers and content providers, either through relaying or tunneling.

1. Introduction

1.1 Purpose and Background

IP multicasting is designed to serve cases where a single source of data content is to be concurrently streamed to multiple recipients. In these types of cases, multicasting provides resource efficiencies (both for the overall network and the content server) relative to unicasting. These efficiencies are possible because of the avoidance of unnecessary duplication of streams, video/audio processing, etc. Multicasting also provides resource efficiencies relative to IP broadcasting in that content data is only delivered to end hosts which request it.

There are many real-life situations where IP multicasting is selected as the technology used for concurrent content delivery of identical content data to multiple end-hosts.   "Requirements for Accounting, Authentication and Authorization in Well Managed IP Multicasting Services", (draft-ietf-mboned-maccnt-req-04.txt, hereafter MACCNT-REQ-draft) describes the requirements in CDN services using IP multicast[1]. "Issues Related to Receiver Access Control in the Current Multicast Protocols" (draft-ietf-mboned-rac-issues-03.txt, hereafter RAC-ISSUES-draft) discusses the requirements and existing support for large-scale, multi-entity content delivery services[2]. The requirements are derived from:

- need for user tracking and billing capabilities
- need for network access control and/or content access control to satisfy the requirements of the CP
- methods for sharing information between the network service provider and content provider to make it possible to fulfill the above two requirements.

Detailed requirements are presented in MACCNT-REQ-draft.   These requirements include mechanisms for recording end-user requests and provider responses for content-delivery, sharing user information (possibly anonymously depending on the trust model) between content provider and network service provider, and protecting resources through the prevention of network and content access by unauthorized users, as well as other AAA related requirements.

The purpose of this memo is to provide a generalized framework for solution standards to meet these requirements. This framework is to provide a basis for defining protocols, but definition of the actual protocols is outside of the scope of this memo.

2. Definitions and Abbreviations

2.1 Definitions

For the purposes of this memo 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

CDN: Content Delivery Network

CDS: Content Delivery Services

CP: Content Provider

NSP: Network Service Provider

TP: Transit Provider

3. Framework and Roles of Entities

3.1 **Framework for multicast AAA allowing bandwidth Management**

A general high-level framework can be represented as follows.

```
            +------------------------------+
            |    user                      |
            |                              |
            +------------------------------+
                | Access       ^ Response
                | Request      | & Multicast Data
                V              |
            +------------------------------+
            |    NSP                       |
            |                              |
            +------------------------------+
                | Access         ^ Response
                | Request        | (Success)
                v                |
            +------------------------------+
            |    CP                        |
            |                              |
            +------------------------------+
```

For the sake of simplicity, the above diagram portrays a case where there is a single NSP entity and a single CP entity. Under the framework it is possible for there to be multiple CPs connected to the same NSP. It is also possible for the same CP to be connected to multiple NSP networks (e.g. network selection). In other words the relationship of NSP:CP can be described as 1:1, 1:N or M:N. Furthermore it is possible that the NSP and CP could be the same entity.

Description of Roles:

The user selects a CP and a NSP when the user requests content. The NSP may be automatically selected by a user terminal: e.g. a fixed line NSP for STB or a mobile NSP for mobile phone.

The CP is responsible for Authentication and Authorization of users' access to content that the CP manages. The CP hopes to collect accounting information related to the access of their content. The CP may choose to authenticate and authorize NSPs which are eligible to provide users access to its contents. When the CP cannot or decides not to provide content to be multicast to users, the CP is responsible for notifying the NSP of the reason.

The NSP is responsible for managing its network resources. The NSP may perform admission control to protect bandwidth resource and needs

authorized information regarding user access for bandwidth management.
It is also responsible for confirming (authentication by proxy) with the CP whether the user is eligible to receive content. When the NSP cannot or decides not to multicast to users, the NSP is responsible for notifying the users of the reason.
In addition to the three basic entities of user, NSP and CP, this AAA framework for multicasting supports transit provision which transfers multicast streams from the CP to the NSP.

3.2 Multiple User IDs

Users may be assigned separate user IDs for each subscription for various NSPs and CPs. When the user wants to access content or otherwise use the network, the user registers the corresponding user ID with a request for content, etc: web authentication is one possible method.

Terminal portability can be realized if the NSP authenticates a user using a user ID. This allows the user to access the content from various network access points.

Each CP may identify users by the user IDs it has issued to them.

The NSP and CP do not need to know the corresponding user id for the same user in the other provider's domain, and it is not necessary that there is a one to one relationship. It is quite possible for one person to hold multiple user ids for the same provider.

3.3 Accounting

The NSP should not manage multicast states on a subnet basis, but on a user basis because the NSP needs to notify start and stop times for accounting purposes. This means that the NSP sends an indication for Join and Leave on a user basis.

The NSP should log both user and host information for each join and leave, indicating the corresponding multicast source for each action. It is important that such log use a standard format so that it can be shared with the CP. Intermittent logs between the join and leave also could serve useful in billing discrepancies, and disconnects without leaves. Ideally a solution would also provide standard ways for the NSP to share logged information with the CP. When shared it is important that the CP be able to match the user to the user within its domain.

3.4 Access Control and CP selection by NSP

When a NSP receives an access request from a user, it is necessary for the NSP to determine to which CP the request is directed. It is necessary for the NSP to ensure that it is not spoofed by an inappropriate CP.

3.5 Network Resource Management by NSP

After authorizing a user request, the NSP may conduct admission control based on its bandwidth management policy. For example, if the NSP manages the shared bandwidth of access lines, the NSP might calculate available bandwidth and necessary bandwidth, and based on these calculations determine to accept or reject a user request.

3.6 Access Control and Distinguishing of Users by CP

The user ID and authentication information are forwarded transparently by the NSP so that the CP can distinguish the user, as well as authenticate and authorize the request.

3.7 Caching of AAA results

An NSP should be able to cache AAA results based an understanding between the NSP and a CP. The AAA cache would store information about permissions of a specific user to receive multicast data from specified channel(s) up to specified expiration date(s) and time(s). If such caching is implemented, a method must exist for the CP to communicate this permission information to the NSP. The NSP refers to the AAA cache and if the cache indicates that the user has permission to receive multicast data from a specific channel at that time, the NSP may forward the data without querying the CP.

It should be possible for a CP to send a directive to the NSP to refresh or change permissions for a user for specific channel(s).

It is necessary for the NSP to requery the CP for authorization should a user be receiving content when the permission expires.

It would be desirable to have a mechanism by which CPs could proactively push permission information to the cache even when not specifically queried by the NSP.

4. Network Connection Model and Functional Components

Section 3.1 introduces the high-level AAA framework for multicasting. This section provides more detail on the network connection model and constituent functional components.

4.1 Basic Connection Model

```
               +-------------+
               |     user    |
               |             |
               +-------------+
                       ^ Access & Resource
                       | Request
                       v
               +-------------+
               |     NSP     |
               |             |
               +-------------+
                       ^ Access
                       | Request
                       v
               +-------------+
               |     CP      |
               |             |
               +-------------+
```

First a user desiring authorization sends an Access request to an NSP which then forwards it on to the appropriate CP for Authentication and Authorization. The CP responds with either "success" or "failure". If "success", the NSP may forward a success response and stream multicast data to the user.

In this model the user selects the NSP to which to send its content request. Based on this request the NSP selects an appropriate CP to which it forwards the request. The CP responds to the NSP's request: it may not respond to another NSP in regards to the request.

In this model, as described in section 3.1, the relationship between NSP and CP can be 1:1, 1:N or M:N. Users may connect to multiple networks, and networks have multiple users.

4.2 Transit Provision

The diagram below shows that a Transit Provider(hereafter, TP) may relay requests between NSPs and CPs.

```
               +-------------+
               |     user    |
               |             |
               +-------------+
                       ^ Access & Resource
                       | Request
                       v
               +-------------+
               |     NSP     |
               |             |
               +-------------+
                       ^ Access & Resource
                       | Request
                       v
               +-------------+
               |     TP      |
               |             |
               +-------------+
                       ^ Access
                       | Request
                       v
               +-------------+
               |     CP      |
               |             |
               +-------------+
```

For the sake of simplification the above diagram shows a 1-1 relationship between an NSP and a TP. However it is also possible for a single NSP to connect to multiple TPs, and a single TP to multiple NSPs.

A single TP may connect to one or more CPs. Similarly just as a single CP may connect to multiple NSPs (as described in the general high-level framework, section 3.1), a single CP may connect to one or more TPs.

A solution will include a mechanism through which the NSPs know which TP(s) are to be used to communicate with which CP(s), and CPs know which TP(s) to use for which NSP(s). When a TP receives an access or resource request from an NSP or CP, it must relay the request to the correct CP or NSP, respectively. Minimally, this means that it must reconstruct the request with translated address information. In this model therefore a TP must understand the format and meaning of the requests.

There may be multiple TPs between a NSP and CP so that a TP is actually receiving from and/or sending requests to another TP and not directly from/to a NSP or CP.

4.3 Transit with Tunnels

In addition to the above model of request relaying, a TP may communicate requests through tunneling based on the contract between the TP and the NSP and/or between the TP and the CP. So in this case the TP will not directly need to process the contents of the access and resource request (such as, header information), but instead pass the request directly to the correct NSP or CP, using a separate protocol to wrap the original requests.

Below is a diagram, representing how a TP can provider tunneling between NSP(s) and CP(s).

```
               +-----------------+
               |     user        |
               |                 |
               +-----------------+
                       ^ Access & Resource
                       | Request
                       v
               +------------------+
               |       NSP        |
               |                  |
               +------------------+
                 |^|
                 |:|
                 |:|
               +-|:|--------------+
               | |:|   TP         |
               | |:|              |
               +-|:|--------------+
                 |:|
                 |:| Tunnel
                 |:|
                 |V|
               +------------------+
               |       CP         |
               |                  |
               +------------------+
```

In this model too, the relationship between NSP and TP and between transit provider and CP can be 1:1, 1:N or M:N.

4.4 Constituent Logical Functional Components of the fully enabled AAA Framework

Section 3.1 introduces the high-level AAA framework for multicasting. Below is a diagram of a fully enabled multicasting network with AAA, including the logical components within the various entities.

```
            +-------------------------------+
            |            user               |
            |                               |
            +-------------------------------+
                          ^
                          | Access & resource request
                          v
            +-------------------------------+
            |            NSP                |
            |                               |
            |+--------------+    +---------+|
            ||NR Management |<-->|AAA Proxy||    (NR= network resource)
            |+--------------+ RR +---------+|    (RR= resource request)
            +-------------------------------+
                          ^
                          | Access request
                          v
            +------------------------------+
            |             CP               |
            |                              |
            |         +---------+          |
            |         |   AAA   |          |
            |         +---------+          |
            +------------------------------+
```

In the fully enabled model the NSP provides proxying of authentication and authorization between the NSP and CP, as well as user-based accounting. The AAA proxy server of the NSP communicates with the CP's AAA server. Although not show in the above diagram for the sake of simplicity, in addition to direct proxying between a NSP and CP, this proxying may be done through a TP. This means that the transit provider too is able to support AAA proxying.

In the fully enabled model the NSP also includes a component that provides network resource management (e.g. QoS management), as described in section 3.4, "Network Resource Management by NSP". When a transit provider is used it may also provide Network Resource management of its own resources.

4.5 Modularity of the framework

In the interest of flexibility, this framework is modular so that it is possible that partially enabled versions of the models are supported. A AAA-enabled version provides AAA functionality without Network Resource management. A Network-Resource-Management-enabled (QoS-enabled) version provides Network Resource management without AAA functionality. Similarly, the possibility of one or more layers of transit provision between an NSP and CP is in the interest of modularity and extendibility.

5. IANA considerations

This memo does not raise any IANA consideration issues.

6. Security considerations

Refer to section 3.3. Also the user information related to authentication with the CP should be protected in some way. Otherwise, this memo does not raise any new security issues which are not already existing in the original protocols. Enhancement of multicast access control capabilities may enhance security performance.

7. Conclusion

This memo provides a generalized framework for solution standards to meet the requirements presented in MACCNT-REQ-draft. Further work should be done to break down the content provider and network service provider entities into their functional objects such as edge devices, AAA servers, etc.

Normative References

[1] Hayashi, et. al., "Accounting, Authentication and Authorization Issues in Well Managed IP Multicasting Services", draft-ietf-mboned-maccnt-req-04.txt, February 2006, Work in Progress.

[2] Hayashi, et. al., "Issues Related to Receiver Access Control in the Current Multicast Protocols", draft-ietf-mboned-rac-issues-03.txt, April 2006, Work in Progress.

Authors' Addresses

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

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

Haixiang He
Nortel
600 Technology Park Drive
Billerica, MA 01801, USA
Phone: +1 978 288 7482
Email: haixiang@nortel.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 December 25, 2006.

Acknowledgement

Funding for the RFC Editor function is currently provided by the

Internet Society.