          **Overview of the Internet Multicast Addressing Architecture**
                    **draft-ietf-mboned-addrarch-00.txt**

Status of this Memo

   This document is an Internet-Draft and is subject to all provisions
   of section 3 of RFC 3667.  By submitting this Internet-Draft, each
   author represents that any applicable patent or other IPR claims of
   which he or she is aware have been or will be disclosed, and any of
   which he or she become aware will be disclosed, in accordance with
   RFC 3668.

   Internet-Drafts are working documents of the Internet Engineering
   Task Force (IETF), its areas, and its working groups.  Note that
   other groups may also distribute working documents as
   Internet-Drafts.

   Internet-Drafts are draft documents valid for a maximum of six months
   and may be updated, replaced, or obsoleted by other documents at any
   time.  It is inappropriate to use Internet-Drafts as reference
   material or to cite them other than as "work in progress."

   The list of current Internet-Drafts can be accessed at
   http://www.ietf.org/ietf/1id-abstracts.txt.

   The list of Internet-Draft Shadow Directories can be accessed at
   http://www.ietf.org/shadow.html.

   This Internet-Draft will expire on May 30, 2005.

Copyright Notice

Abstract

   The lack of up-to-date documentation on IP multicast address
   allocation and assignment procedures has caused a great deal of
   confusion.  To clarify the situation, this memo describes the
   allocation and assignment techniques and mechanisms currently (as of
   this writing) in use.

Table of Contents

## 1. Introduction

Good, up-to-date documentation of IP multicast is close to non-existent.  Particularly, this is an issue with multicast address allocations (to networks and sites) and assignments (to hosts and applications).  This problem is stressed by the fact that there exists confusing or misleading documentation on the subject [RFC2908].  The consequence is that those who wish to learn of IP multicast and how the addressing works do not get a clear view of the current situation.

The aim of this document is to provide a brief overview of multicast addressing and allocation techniques.  The term 'addressing architecture' refers to the set of addressing mechanisms and methods in an informal manner.

It is important to note that Source-specific Multicast (SSM) [I-D.ietf-ssm-arch] does not have these addressing problems; hence, this document focuses on Any Source Multicast (ASM) model.  The applicability of SSM has been briefly discussed in [I-D.ietf-mboned-ipv6-multicast-issues].

This memo obsoletes RFC 2908 and RFC 2909 and re-classifies them Historic.

### 1.1  Terminology: Allocation or Assignment

Almost all multicast documents and many other RFCs (such as DHCPv4 [RFC2131] and DHCPv6 [RFC3315]) have used the terms address "allocation" and "assignment" interchangeably.  However, the operator and address management communities use these for two conceptually different processes.

In unicast operations, address allocations refer to leasing a large block of addresses from Internet Assigned Numbers Authority (IANA) to a Regional Internet Registry (RIR), from RIR to a Local Internet Registry (LIR) possibly through a National Internet Registry (NIR). Address assignments, on the other hand, are the leases of smaller address blocks or even single addresses to the end-user sites or end-users themselves.

Therefore, in this memo, we will separate the two different functions: "allocation" describes how larger blocks of addresses are obtained by the network operators, and "assignment" describes how applications, nodes or sets of nodes obtain a multicast address for their use.

[[ NOTE IN DRAFT: is this choice of terminology too confusing? ]]

## 2. Multicast Address Allocation

Multicast address allocation, i.e., how a network operator might be able to obtain a larger block of addresses, can be handled in a number of ways as described below.

Note that these are all only pertinent to ASM -- SSM requires no address block allocation because the group address has only local significance (however, the address assignment inside the node is still an issue discussed in Section 3.2).

### 2.1 Derived Allocation

Derived allocations take the unicast prefix or some other properties of the network to determine unique multicast address allocations.

### 2.1.1 GLOP Allocation

GLOP address allocation [RFC3180] inserts the 16-bit public Autonomous System (AS) number in the middle of the IPv4 multicast prefix 233.0.0.0/8, so that each AS number can get a /24 worth of multicast addresses.  While this is sufficient for multicast testing or small scale use, it might not be sufficient in all cases for extensive multicast use.

A minor operational debugging issue with GLOP addresses is that the connection between the AS and the prefix is not apparent from the prefix, but has to be calculated (e.g., from [RFC3180], AS 5662 maps to 233.22.30.0/24).  A usage issue is that GLOP addresses are not tied to any prefix but to routing domains, so they cannot be used or calculated automatically.

### 2.1.2 Unicast-prefix -based Allocation

RFC 3306 [RFC3306] describes a mechanism which embeds up to 64 first bits of an IPv6 unicast address in the prefix part of the IPv6 multicast address, leaving at least 32 bits of group-id space available after the prefix mapping.

A similar mapping has been proposed for IPv4 [I-D.ietf-mboned-ipv4-uni-based-mcast], but it provides a rather low amount of addresses (e.g., 1 per an IPv4 /24 block).  While there exist large networks without an AS number of their own, this has not been seen to add sufficient value compared to GLOP addressing.

The IPv6 unicast-prefix -based allocations are an extremely useful way to allow each network operator, even each subnet, obtain multicast addresses easily, through an easy computation.  Further, as

the IPv6 multicast header also includes the scope value [RFC3513], multicast groups of smaller scope can also be used with the same mapping.

The IPv6 Embedded RP technique [RFC3956], used with Protocol Independent Multicast - Sparse Mode (PIM-SM), further leverages the unicast prefix based allocations, by embedding the unicast prefix and interface identifier of the PIM-SM Rendezvous Point (RP) in the prefix.  This provides all the necessary information needed to the routing systems to run the group in either inter- or intra-domain operation.  A difference to RFC 3306 is, however, that the hosts cannot calculate their "multicast prefix" automatically, as the prefix depends on the decisions of the operator setting up the RP but rather needs to be communicated somehow.

All the IPv6 unicast-prefix -based allocation techniques provide sufficient amount of multicast address space for the network operators.

## 2.2  Scope-relative Allocation

Administratively scoped multicast [RFC2365] is provided by two different means: under 239.0.0.0/8 in IPv4 or by 4-bit encoding in the IPv6 multicast address prefix [RFC3513].

As IPv6 scope-relative allocations can be handled with unicast-prefix -based multicast addressing as described in Section 2.1.2, and there is no need for separate scope-relative allocations, we'll just discuss IPv4 in this section.

The IPv4 scope-relative prefix 239.0.0.0/8 is further divided to Local Scope (239.255.0.0/16) and Organization Local Scope (239.192.0.0/14); other parts of the administrative scopes are either reserved for expansion or undefined [RFC2365].

Topologies which act under a single administration can easily use the scoped multicast addresses for their internal groups.  Groups which need to be shared between multiple routing domains (but not propagated through Internet) are more problematic and typically need an assignment of a global multicast address because their scope is undefined.

There is a large number of multicast applications (such as "Norton Ghost") which are restricted either to a link or a site, but it is extremely undesirable to propagate them further (either to the rest of the site, or beyond the site).  Typically many such applications have been given a static IANA address assignment; this makes it challenging to implement proper propagation limiting -- which could

be easier if such applications could have been assigned specific
scope-relative addresses instead.  This is an area of further future
work -- it might be able to mitigate this issue if there was more
coordination inside the scope-relative allocation block.

## 2.3  Static IANA Allocation

In some rare cases, some organizations may have been able to obtain
static multicast address allocations directly from IANA.  Typically
these have been meant as a block of static assignments to multicast
applications, as described in Section 3.4.  In principle, IANA does
not allocate multicast address blocks to the operators but GLOP or
Unicast-prefix -based allocations should be used instead.

## 2.4  Dynamic Allocation

RFC 2908 [RFC2908] proposed three different layers of multicast
address allocation and assignment, where layers 3 (inter-domain
allocation) and layer 2 (intra-domain allocation) could be applicable
here.  Multicast Address-Set Claim Protocol (MASC) [RFC2909] is an
example of the former, and Multicast Address Allocation Protocol
(AAP) [I-D.ietf-malloc-aap] (abandoned in 2000 due lack of interest
and technical problems) is an example of the latter.

Both of the proposed allocation protocols were quite complex, and
have never been deployed or seriously implemented.

It can be concluded that there are no dynamic multicast address
allocation protocols, and other methods such as GLOP or
unicast-prefix -based addressing should be used instead.

## 3.  Multicast Address Assignment

For multicast address assignment, i.e., how an application learns the
address it can use, or a node (or a set of nodes) learns an address
it could use for an application, has a number of options as described
below.

Any IPv6 address assignment method should be aware of the guidelines
for the assignment of the group-IDs for IPv6 multicast addresses
[RFC3307].

## 3.1  Derived Assignment

There are significantly fewer options for derived address assignment
compared to derived allocation.  Derived multicast assignment is only
being specified for IPv6 link-scoped multicast
[I-D.ietf-ipv6-link-scoped-mcast], where the EUI64 is embedded in the

multicast address, providing a node with unique multicast addresses
for link-local ASM communications.

### 3.2  SSM Assignment inside the Node

While the SSM multicast addresses have only local (to the node)
significance, there is still a minor issue on how to assign the
addresses between the applications running on the same node (or more
precisely, an IP address).

This assignment is not considered to be a problem because typically
the addresses for the applications are selected manually or
statically, but if done using an API, the API could check that the
addresses do not conflict prior to assigning one.

### 3.3  Manually Configured Assignment

With manually configured assignment, the network operator which has a
multicast address prefix assigns the multicast group addresses to the
requesting nodes using a manual process.

Typically the user or administrator which wants to use a multicast
address for particular application requests an address from the
network operator using phone, email, or similar means, and the
network operator provides the user with a multicast address.  Then
the user/administrator of the node or application manually configures
the application to use the assigned multicast address.

This is a relatively simple process in the beginning, but would
become unscalable if the multicast usage would get on a serious rise
(fortunately, we have dynamic assignment, see Section 3.5).  Another,
separate issue is to ensure that the users wishing to use that
application are able to locate the configured multicast address
("rendezvous" or "service discovery"); in this particular case, this
might call for e.g., DNS-based discovery of the multicast address.

This is the most commonly used technique when the multicast
application does not have a static IANA assignment.

### 3.4  Static IANA Assignment

In contrast to manually configured assignment, as described above,
static IANA assignment refers to getting a globally unique assignment
for the particular application directly from IANA.  Guidelines for
IANA are described in [I-D.ietf-mboned-rfc3171bis].

This is seen as lucrative because it's the simplest approach for
application developers because they can then hard-code the multicast

address, requiring no lease of the usable multicast address, and
likewise the client applications do not need to perform any kind of
service discovery (but depending on hard-coded addresses).  However,
this is a bad approach architecturally, as we should focus on
enhancing and deploying service discovery and address assignment (as
needed) instead of encouraging a "land-grab" of multicast addresses.

In summary, there are applications which have obtained a static IANA
assignment, some of which are really needed, and some of which
probably should not have been granted.

## 3.5  Dynamic Assignments

The layer 1 of RFC 2908 [RFC2908] described dynamic assignment from
Multicast Address Allocation Servers (MAAS) to applications and
nodes, with Multicast Address Dynamic Client Allocation Protocol
(MADCAP) [RFC2730] and a proposal for DHCPv6 assignments
[I-D.jdurand-assign-addr-ipv6-multicast-dhcpv6] as examples.

Based on a multicast prefix, it would be rather straightforward to
deploy a dynamic assignment protocol which would lease group
addresses to the applications wishing to use multicast.  For example,
only few have implemented MADCAP, and it's not significantly
deployed.  Moreover, it is not clear how widely for example the APIs
for communication between the multicast application and the MADCAP
client operating at the host have been implemented [RFC2771].

Based on that, a conclusion is that multicast is that either:

1.  multicast is not significantly attractive in the first place,

2.  manually configured assignments are sufficient for now, or

3.  that there are other gaps why dynamic assignments are not seen as
    a useful approach (for example, issues related to service
    discovery/rendezvous).

In consequence, more work on rendezvous/service discovery will be
needed to make dynamic assignment more useful.

## 3.6  Future Developments

IPv6 could offer an alternative to dynamic assignments due to its
larger address space: if a multicast prefix (e.g., about 2^32 bits
worth of group-id's) is allocated to a subnet, it would be sufficient
to ensure that multiple applications running on that subnet do not
try to use the same address (selected e.g., using a random process).
This could call for a Duplicate Address Detection process, and/or a

way for the RPs to inform the hosts about the prefix that could be
used on each subnet (assuming Embedded-RP would be used).

4.  **Multicast Address Discovery**

[[ NOTE IN DRAFT: it is not clear whether this section belongs in
this document at all; it is somewhat related, but could bear a more
extensive discussion elsewhere.  It should likely go in a separate
document (if there was one discussing these problems!), or in an
appendix.  Feedback is appreciated.  ]]

As was noted in Section 3, multicast address discovery (i.e., service
discovery or "rendezvous") is a problem with multicast address
assignment.  In particular, an acceptable mechanism (mechanisms such
as Service Location Protocol (SLP) [RFC2608] seem to have been
considered too complex) seems to be missing which the hosts wishing
to participate in a group could use to find the address of that group
[MBONED-IETF59].

It is worth noting that as long as not deploying an address
assignment and service discovery protocols/mechanisms means that one
can get a static address assignment from IANA, there is little
interest from the application developers to actually do anything
except try to get the assignment from IANA.  Conclusion: if we want
to use non-IANA processes, the assignments must be either forbidden
completely, or made sufficiently difficult that it's easier for the
application developers to take another route if a feasible mechanism
is available.

There are two issues in the service discovery:

1.  The session initiator being able to publish the session somehow,
    and

2.  The session participants finding out about the session (rather
    than creating their own).

When manually configured or static IANA assignments are used, 1)
should be relatively straightforward (if something needs to be
manually configured or statically assigned, putting it e.g., in DNS
should not be a problem).  However, this is still more complex for
dynamic or derived assignments because it implies that the host or
the application has the right to make that publication on its own,
rather than through a manual process by an administrator.

2) is always a challenge, but could leverage for example DNS (e.g.,
by relying on using SRV records with the DNS search path, as
described in [I-D.iab-dns-choices] and

[I-D.palet-v6ops-tun-auto-disc]).

## 5.  Summary and Future Directions

This section summarizes the mechanisms and analysis discussed in this
memo, and presents some potential future directions.

## 5.1  Prefix Allocation

Summary of prefix allocation methods is in Figure 1.

```
+-------+-------------------------------+--------+--------+
| Sect. | Prefix allocation method      | IPv4   | IPv6   |
+-------+-------------------------------+--------+--------+
|   2   | SSM                           | NoNeed | NoNeed |
| 2.1.1 | Derived: GLOP                 |  Yes   | NoNeed*|
| 2.1.2 | Derived: Unicast-prefix -based|No -yet |  Yes   |
|  2.2  | Separate Scope-relative       |  Yes   | NoNeed*|
|  2.3  | Static IANA allocation        |  No    |  No    |
|  2.4  | Dynamic allocation protocols  |  No    |  No    |
+-------+-------------------------------+--------+--------+
```
* = the need satisfied by IPv6 unicast-prefix -based allocation.

Figure 1

o  Only ASM is affected by the assignment/allocation issues (however,
   both ASM and SSM have roughly the same address discovery issues).

o  GLOP allocations seem to provide a sufficient IPv4 multicast
   allocation mechanism for now, but could be extended in future.
   Scope-relative allocations provide the opportunity for internal
   IPv4 allocations.

o  Unicast-prefix -based addresses and the derivatives provide good
   allocation strategy with IPv6, also for scoped multicast
   addresses.

o  Dynamic allocations are a too complex and unnecessary mechanism.

o  Static IANA allocations are an architecturally unacceptable
   approach.

## 5.2  Address Assignment

Summary of address assignment methods is in Figure 2.

```
+-------+-------------------------------+----------+----------+
| Sect. | Address assignment method     | IPv4     | IPv6     |
+-------+-------------------------------+----------+----------+
|  3.1  | Derived: link-scope addresses |   No     |   Yes    |
|  3.2  | SSM (inside the node)         |   Yes    |   Yes    |
|  3.3  | Manual assignment             |   Yes    |   Yes    |
|  3.4  | Static IANA assignment        |LastResort|LastResort|
|  3.5  | Dynamic assignment protocols  |   Yes    |   Yes    |
+-------+-------------------------------+----------+----------+
```

Figure 2

o  Manually configured assignment is what's typically done today, and
   works to a sufficient degree in smaller scale.

o  Static IANA assignment has been done extensively in the past, but
   it needs to be tightened down to prevent problems caused by
   "land-grabbing".

o  Dynamic assignment, e.g., using MADCAP have been implemented, but
   there is no wide deployment, so a solution is there -- but either
   there are other gaps in the multicast architecture or there is no
   need for it in the first place, when manual configuration is
   possible, and static IANA assignments are still there.

o  Derived assignments are only applicable in a fringe case of
   link-scoped multicast.


## 5.3  Future Actions

o  Multicast address discovery/"rendezvous" needs to be analyzed at
   more length, and an adequate solution provided; the result also
   needs to be written down to be shown to the IANA static assignment
   requestors.

o  IPv6 multicast DAD and/or multicast prefix communication
   mechanisms should be analyzed: whether there is demand or not, and
   specify if so.

o  The IETF should consider whether to specify more ranges of the
   IPv4 scope-relative address space for static allocation for
   applications which should not be routed over the Internet (such as
   backup software, etc.  -- so that these wouldn't need to use

global addresses which should never leak in any case).

o  The IETF should seriously consider its static IANA allocations
   policy, e.g., "locking it down" to a stricter policy (like "IETF
   Consensus") and looking at developing the discovery/rendezvous
   functions, if necessary.


## 6.  Acknowledgements

Tutoring a couple multicast-related papers, the latest by Kaarle
Ritvanen [RITVANEN] convinced the author that the up-to-date
multicast address assignment/allocation documentation is necessary.

Multicast address allocations/assignments were discussed at the
MBONED WG session at IETF59 [MBONED-IETF59].

Dave Thaler, James Lingard, and Beau Williamson provided useful
feedback for the preliminary version of this memo.  Myung-Ki Shin
also suggested improvements.

## 7.  IANA Considerations

This memo includes no request to IANA, but as the allocation and
assignment of multicast addresses are related to IANA functions, it
wouldn't hurt if the IANA reviewed this entire memo.

IANA considerations in sections 4.1.1 and 4.1.2 of [RFC2908] still
apply to the administratively scoped prefixes.

(RFC-editor: please remove this section at publication.)

## 8.  Security Considerations

This memo only describes different approaches to allocating and
assigning multicast addresses, and this has no security
considerations; the security analysis of the mentioned protocols is
out of scope of this memo.

Obviously, especially the dynamic assignment protocols are inherently
vulnerable to resource exhaustion attacks, as discussed e.g., in
[RFC2730].

## 9.  References

## 9.1  Normative References

[I-D.ietf-ipv6-link-scoped-mcast]

                 Park, J., Shin, M. and H. Kim, "Link Scoped IPv6 Multicast
                 Addresses", draft-ietf-ipv6-link-scoped-mcast-06 (work in
                 progress), October 2004.

   [I-D.ietf-mboned-rfc3171bis]
                 Albanna, Z., Almeroth, K., Cotton, M. and D. Meyer, "IANA
                 Guidelines for IPv4 Multicast Address Assignments",
                 draft-ietf-mboned-rfc3171bis-02 (work in progress), March
                 2004.

   [I-D.ietf-ssm-arch]
                 Holbrook, H. and B. Cain, "Source-Specific Multicast for
                 IP", draft-ietf-ssm-arch-06 (work in progress), September
                 2004.

   [RFC2365]    Meyer, D., "Administratively Scoped IP Multicast", BCP 23,
                 RFC 2365, July 1998.

   [RFC3180]    Meyer, D. and P. Lothberg, "GLOP Addressing in 233/8", BCP
                 53, RFC 3180, September 2001.

   [RFC3306]    Haberman, B. and D. Thaler, "Unicast-Prefix-based IPv6
                 Multicast Addresses", RFC 3306, August 2002.

   [RFC3307]    Haberman, B., "Allocation Guidelines for IPv6 Multicast
                 Addresses", RFC 3307, August 2002.

   [RFC3513]    Hinden, R. and S. Deering, "Internet Protocol Version 6
                 (IPv6) Addressing Architecture", RFC 3513, April 2003.

   [RFC3956]    Savola, P. and B. Haberman, "Embedding the Rendezvous
                 Point (RP) Address in an IPv6 Multicast Address", RFC
                 3956, November 2004.

9.2  **Informative References**

   [I-D.iab-dns-choices]
                 Faltstrom, P. and R. Austein, "Design Choices When
                 Expanding DNS", draft-iab-dns-choices-00 (work in
                 progress), October 2004.

   [I-D.ietf-malloc-aap]
                 Handley, M. and S. Hanna, "Multicast Address Allocation
                 Protocol (AAP)", June 2000.

   [I-D.ietf-mboned-ipv4-uni-based-mcast]
                 Thaler, D., "Unicast-Prefix-based IPv4 Multicast
                 Addresses", draft-ietf-mboned-ipv4-uni-based-mcast-02

                    (work in progress), October 2004.

   [I-D.ietf-mboned-ipv6-multicast-issues]
                    Savola, P., "IPv6 Multicast Deployment Issues",
                    draft-ietf-mboned-ipv6-multicast-issues-01 (work in
                    progress), September 2004.

   [I-D.jdurand-assign-addr-ipv6-multicast-dhcpv6]
                    Durand, J., "IPv6 multicast address assignment with
                    DHCPv6",
                    draft-jdurand-assign-addr-ipv6-multicast-dhcpv6-00 (work
                    in progress), June 2004.

   [I-D.palet-v6ops-tun-auto-disc]
                    Palet, J. and M. Diaz, "Analysis of IPv6 Tunnel End-point
                    Discovery Mechanisms", draft-palet-v6ops-tun-auto-disc-02
                    (work in progress), October 2004.

   [MBONED-IETF59]
                    "MBONED WG session at IETF59",
                    <http://www.ietf.org/proceedings/04mar/172.htm>.

   [RFC2131]  Droms, R., "Dynamic Host Configuration Protocol", RFC
                    2131, March 1997.

   [RFC2608]  Guttman, E., Perkins, C., Veizades, J. and M. Day,
                    "Service Location Protocol, Version 2", RFC 2608, June
                    1999.

   [RFC2730]  Hanna, S., Patel, B. and M. Shah, "Multicast Address
                    Dynamic Client Allocation Protocol (MADCAP)", RFC 2730,
                    December 1999.

   [RFC2771]  Finlayson, R., "An Abstract API for Multicast Address
                    Allocation", RFC 2771, February 2000.

   [RFC2908]  Thaler, D., Handley, M. and D. Estrin, "The Internet
                    Multicast Address Allocation Architecture", RFC 2908,
                    September 2000.

   [RFC2909]  Radoslavov, P., Estrin, D., Govindan, R., Handley, M.,
                    Kumar, S. and D. Thaler, "The Multicast Address-Set Claim
                    (MASC) Protocol", RFC 2909, September 2000.

   [RFC3315]  Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C. and
                    M. Carney, "Dynamic Host Configuration Protocol for IPv6
                    (DHCPv6)", RFC 3315, July 2003.

   [RITVANEN]
               Ritvanen, K., "Multicast Routing and Addressing", HUT
               Report, Seminar on Internetworking, May 2004,
               <http://www.tml.hut.fi/Studies/T-110.551/2004/papers/>.


Author's Address

   Pekka Savola
   CSC - Scientific Computing Ltd.
   Espoo
   Finland

   EMail: psavola@funet.fi

Appendix A.  Open Issues

   (This section will be removed or merged with the rest before
   publication..)

   o  Is the case for IPv4 Unicast-Prefix Base Multicast addressing
      sufficiently strong, or could those organizations just get an AS
      number themselves if they really wanted to do multicast?

   o  Should one merge the routing architecture document's contents here
      as well?