Network Working Group
Internet-Draft
Expires: September 10, 2009

D. Thaler
Microsoft

March 9, 2009

# Unicast-Prefix-based IPv4 Multicast Addresses
# draft-ietf-mboned-ipv4-uni-based-mcast-06.txt

## Status of this Memo



## Copyright Notice



## Abstract

This specification defines an extension to the multicast addressing architecture of the IP Version 4 protocol. The extension presented in this document allows for unicast-prefix-based assignment of multicast addresses. By delegating multicast addresses at the same

time as unicast prefixes, network operators will be able to identify their multicast addresses without needing to run an inter-domain allocation protocol.

Table of Contents

## 1. Introduction

RFC 3180 [RFC3180] defined an experimental allocation mechanism (called "GLOP") in 233/8 whereby an Autonomous System (AS) number is embedded in the middle 16 bits of an IPv4 multicast address, resulting in 256 multicast addresses per AS. Advantages of this mechanism include the ability to get multicast address space without an inter-domain multicast address allocation protocol, and the ease of determining the AS that was assigned the address for debugging and auditing purposes.

Some disadvantages of GLOP include:

- RFC 4893 [RFC4893] expands the size of an AS number to 4 bytes, and GLOP cannot work with 4-byte AS numbers.
- When an AS covers multiple sites or organizations, administration of the multicast address space within an AS must be handled by other mechanisms, such as manual administrative effort or MADCAP [RFC2730].
- During debugging, identifying the AS does not immediately identify the correct organization when an AS covers multiple organizations.
- Only 256 addresses are automatically available per AS, and obtaining any more requires administrative effort.

More recently, a mechanism [RFC3306] has been developed for IPv6 that provides a multicast range to every IPv6 subnet, which is at a much finer granularity than an AS. As a result, the first three disadvantages above are avoided (and the last disadvantage does not apply to IPv6 due to the extended size of the address space).

Another advantage of providing multicast space to a subnet, rather than just to an entire AS, is that multicast address assignment within the range need only be coordinated within the subnet.

This draft specifies a mechanism similar to [RFC3306], whereby a range of global IPv4 multicast address space is provided to each organization that has unicast address space. A resulting advantage

over GLOP is that the mechanisms in IPv4 and IPv6 become more similar.

## 2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].





## 3. Address Space

(RFC-editor: replace TBD in this section and the next with IANA-assigned value, and delete this note.)

A multicast address with the prefix TBD/8 indicates that the address is a Unicast-Based Multicast (UBM) address. The remaining 24 bits are used as follows:

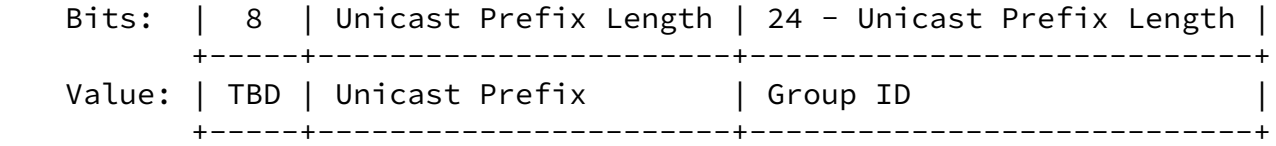


For organizations with a /24 or shorter prefix, the unicast prefix of the organization is appended to the common /8. Any remaining bits may be assigned by any mechanism the organization wishes.

For example, an organization that has a /16 prefix assigned might choose to assign multicast addresses manually from the /24 multicast prefix derived from the above method. Alternatively, the organization might choose to delegate the use of multicast addresses to individual subnets that have a /24 or shorter unicast prefix, or it might choose some other method.

Organizations with a prefix length longer than 24 do not receive any multicast address space from this mechanism; in such cases, another mechanism must be used.

Compared to GLOP, an AS will receive more address space via this mechanism if it has more than a /16 for unicast space. An AS will receive less address space than it does from GLOP if it has less than a /16.

The organization that is assigned the UBM address can be determined by taking the multicast address, shifting it left by 8 bits, and identifying who has been assigned the address space covering the resulting unicast address.

The embedded unicast prefix MUST be a global unicast prefix (i.e., no loopback, multicast, link-local, or private-use IP address space). In addition, since global unicast addresses are not permanently assigned, UBM addresses MUST NOT be hard-coded in applications.

## 4. Examples

The following are a few examples of the structure of unicast-prefix





based multicast addresses.

- Consider an organization that has been assigned the global unicast address space 192.0.2.0/24. This means that organization can use the global multicast address TBD.192.0.2 without coordinating with any other entity. Someone who sees this multicast address and wants to find who is using it can mentally shift the address left by 8 bits to get 192.0.2.0, and then look up who has been assigned unicast address space that includes that address.
- Consider an organization has been assigned a larger address space, x.y.0.0/16. This organization can use the global multicast address space TBD.x.y.0/24 without coordinating with any other entity, and can assign addresses within this space by any mechanism the organization wishes. Someone who sees a multicast address (say) TBD.x.y.10, and wants to find who is using it can mentally shift the address left by 8 bits to get x.y.10.0, and can then look up who has been assigned unicast address space that includes that address.

## 5. Security Considerations

The same well known intra-domain security techniques can be applied

as with GLOP. Furthermore, when dynamic allocation is used within a prefix, the approach described here may have the effect of reduced exposure to denial of space attacks, since the topological area within which nodes compete for addresses within the same prefix is reduced from an entire AS to only within an individual organization or an even smaller area.

## 6. IANA Considerations

IANA should assign a /8 in the global IPv4 multicast address space for this purpose.

## 7. Acknowledgments

This document was updated based on feedback from the MBoneD working group. In particular, Tim Chown, Toerless Eckert, Prashant Jhingran, Peter Koch, John Linn, Dave Meyer, Pekka Savola, Greg Shepherd, and Stig Venaas provided valuable suggestions on the text.

## 8. References





### 8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

### 8.2. Informative References

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

[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.

[RFC4893] Vohra, Q. and E. Chen, "BGP Support for Four-octet AS Number Space", RFC 4893, May 2007.

## Author's Address

Dave Thaler
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
USA

Phone: +1 425 703 8835
Email: dthaler@microsoft.com