

Unicast-Prefix-based IPv4 Multicast Addresses

draft-ietf-mboned-ipv4-uni-based-mcast-03.txt

Dave Thaler

dthaler@microsoft.com

Current Status

- “BGP Support for Four-octet AS Number Space” (draft-ietf-idr-as4bytes-13.txt) now in RFC-editor queue as a Proposed Standard
- Draft-03 submitted with no change from -02 other than updated boilerplate and references
- One open issue on the list

Is space delegated to the subnet, or to the superblock?

- Currently draft implies to subnet, is this too restrictive?
- Comparison with IPv6:
 - RFC 3306 allows *both* by embedding the subnet prefix length in the multicast address, but not enough bits to do so in IPv4

Subnet vs Superblock

- Subnet:
 - Allows for possible allocation procedures/protocols that only require coordination within the subnet
 - For debugging, can easily identify owner down to subnet granularity
 - No ambiguity of granularity
 - More convenient for BIDIR PIM
- Superblock:
 - More flexible allocation within a multi-subnet org
 - Means more usable addresses without constraining the topology
 - Allows use within an org using subnet prefixes longer than /24
- So far feedback has been in favor of superblocks rather than subnets
- Proposal: reword to delegate to (smallest) superblock owner and leave it to superblock owner whether it is delegated to subnets or not