The Multicast Address Allocation Architecture

***Mark Handley**, USC/ISI*

***Dave Thaler**, UMich*

***Deborah Estrin**, USC/ISI*

Goal of this work

To produce an archutecture for multicast address allocation for the global internet.

Requirements:

- Allocate addresses on demand with very high service availablility.
- Minimise probability of address allocation clashes.
- Maximise usage of the limited address space.
- Allocate addresses in a manner that aids scalable multicast routing.

Non Goal

We do not believe it is possible to satisfy these requirements with an architecture that *guarantees* unique address allocation.

- Architectures that guarantee uniqueness fail to deliver on one or both of the requirements of availability or address space packing.

We do not believe it is necessary to guarantee unique allocation.

- We cannot stop people just using addresses without using the allocation architecture.

Architecture Overview

The Architecture Consists of Four parts:

- An (abstract) API that an application can use to request an address.
- A local protocol that a host can use to request an address from a Multicast Address Allocation Server (MAAS).
- A domain-wide announcement protocol that ‘reserves’ individual addresses (on timescales of seconds).
- An inter-domain hierarchical protocol that provides (aggregatable) address sets to domains and provides information to inter-domain multicast routing.

Abstract API

- **Ross Finlayson will present this.**
- **Generally it mirrors what is required from the local address request protocol, but shouldn't be tied to any particular request protocol.**

Requesting an Address from a Local Server

M-DHCP is a candidate for this protocol.

The requirements are:

- for a host to be able to discover the location of a MAAS server,
- to be able to request information about relevant admin scope zones
- to be able to request one or more multicast addresses for a period of time.
- to be able to request an extension of an existing address lifetime.
- to be able to return an address when it is no longer needed.

Domain Wide Address Allocation

AAP is our candidate for this protocol.

The requirements are:

- For many servers within the same domain to be able to allocate addresses from a given address set.
- To achieve near 100% address space utilisation.
- To achieve low probability of address collision.
- To allocate addresses with a maximum delay of a few seconds.
- To use relatively little bandwidth.
- To be incrementally deployable.
- To be always available, and to be very resilient to network failures.

Inter-domain Multicast Address Set Allocation

- **MASC is our candidate for this protocol**
- **Requirements**
  - To allocate address sets for AAP to allocate addresses out of.
  - To do so in a way that has an *extremely* low probability of address set collision.
  - To do so in a way that address sets hierarchically agregate so that they may be useful to scalable inter-domain multicast routing protocols.
  - To be adaptive to long-term changes in topology, but not to short term failures.
  - To operate on timescales of days.

Note on MASC and multicast routing

It's easy to get confused about the relationship between MASC, BGP and BGMP:

- **They run in the same routers, but are conceptually separate.**
  - MASC builds a hierarchy and allocate address sets.
  - BGP+ will then distribute these address sets in a G-RIB.
  - BGMP uses the G-RIB to do multicast routing.
- **MASC is not tied to BGP or BGMP but we expect it to be used this way in practice.**

Allocation Domains

- **We expect that an allocation domain will typically be an AS.**
- **If a unicast AS is too large, it can be spilt into multiple multicast AS's, and BGMP run at the internal border.**
- **If a unicast AS is allocating too few addresses and is singly connected, it should not be its own domain but should be part of its parent domain instead.**
  - This can be configured to happen adaptively.

Example