The malloc Architecture

Steve Hanna

steve.hanna@sun.com

IP Multicast Model

- Group is identified with class D IP address
- Any host can send to a multicast address
- To receive, join the group via IGMP

- Packets can be constrained using TTL or admin-scoped multicast addresses

Hard Parts

- Multicast addresses are relatively scarce
- Multicast routing requires state
  - Finding core for a group
  - Maintaining a distribution tree

- Multicast address allocation can help with both of these

Ideal malloc Requirements

- Collision-Free
- High Address Space Utilization
- Aggregatable Address Allocation
- Robust
- Scalable
- Secure
- Fast and Efficient

The Bad News

- We can’t do all that at the same time
  - See Mark Handley’s thesis
- So we split up the problem space
  - Interdomain, Intradomain, and Host-Server
- And relax a few overall requirements
  - Mostly Collision-Free
  - Decent Address Space Utilization

Three Protocols

- Interdomain: MASC
- Intradomain: AAP
- Host-server: MDHCP

MASC

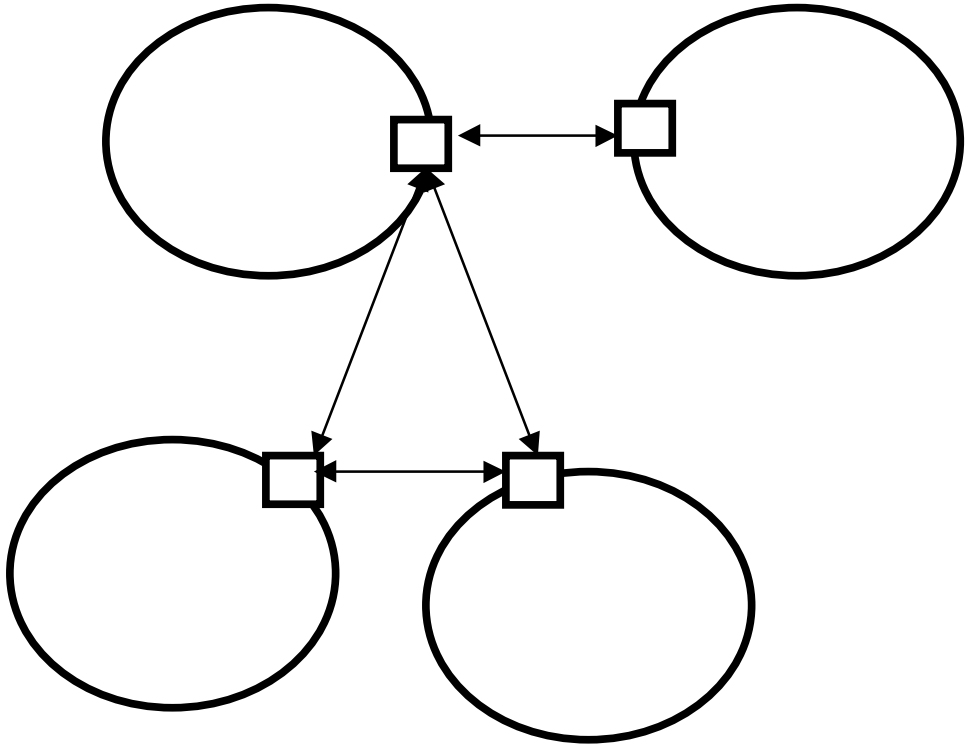

MASC

- Interdomain
- Hierarchy of MASC domains
- Time scale: days (to span outages)
- Parent injects block
- Siblings send out claims and complain if they collide
- Parent can enforce policy

AAP

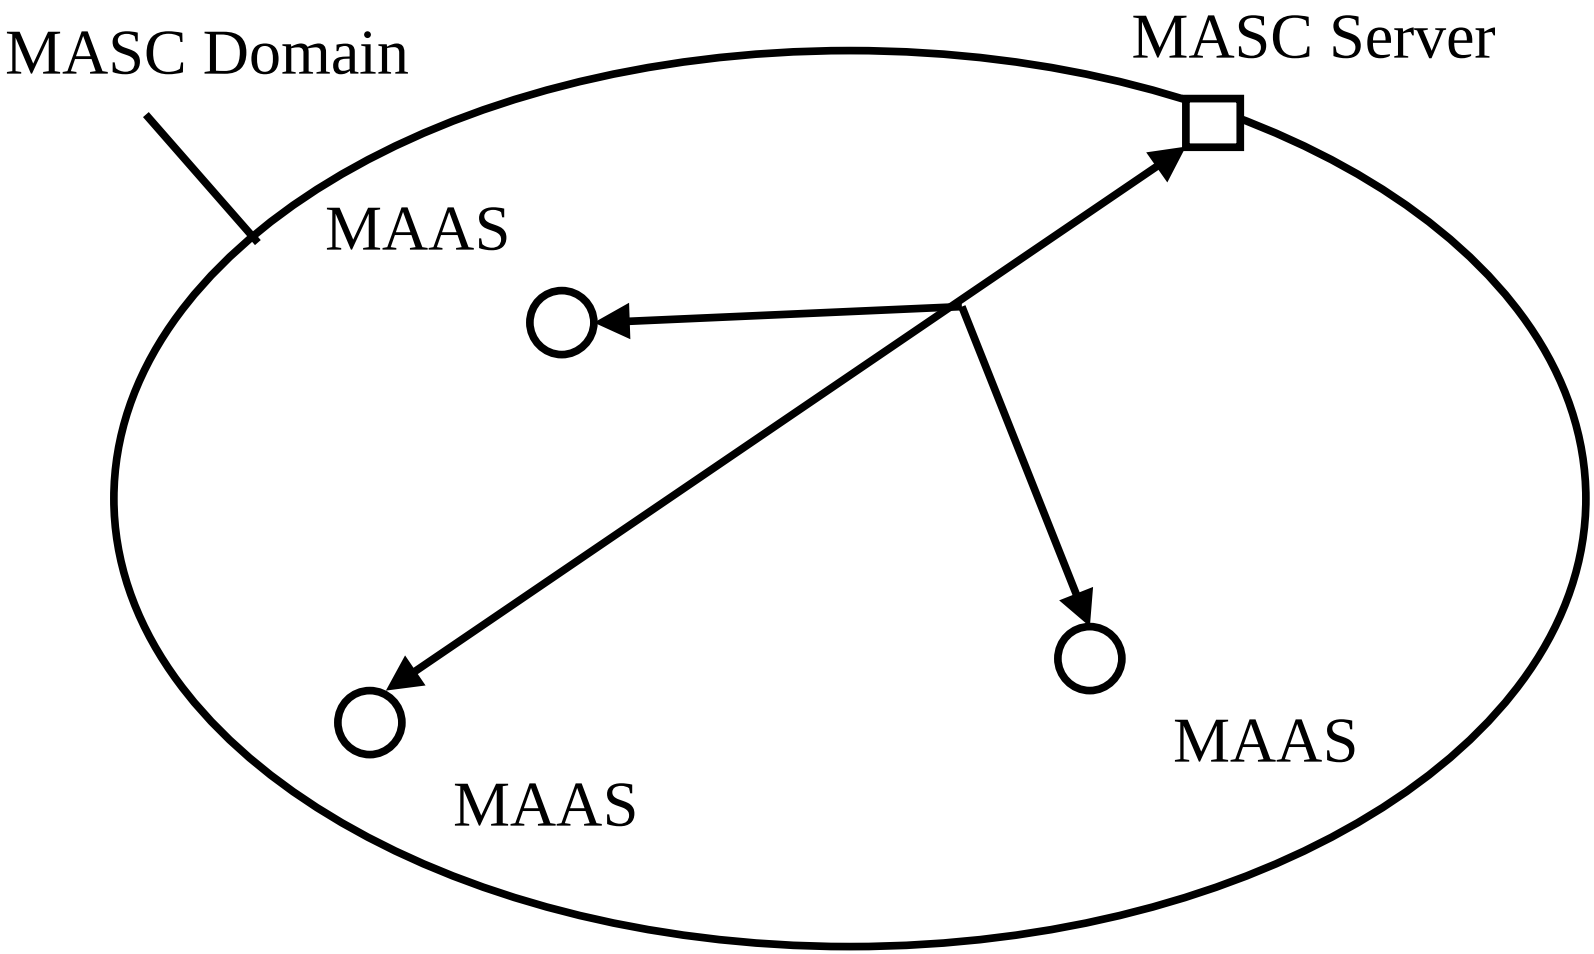

AAP

- Intradomain
- Peer to peer via multicast
- MASC server announces address sets
- AAP servers claim addresses from the sets and announce those claims periodically
- Uses periodic multicast announcements (like SAP)

MDHCP Server Discovery

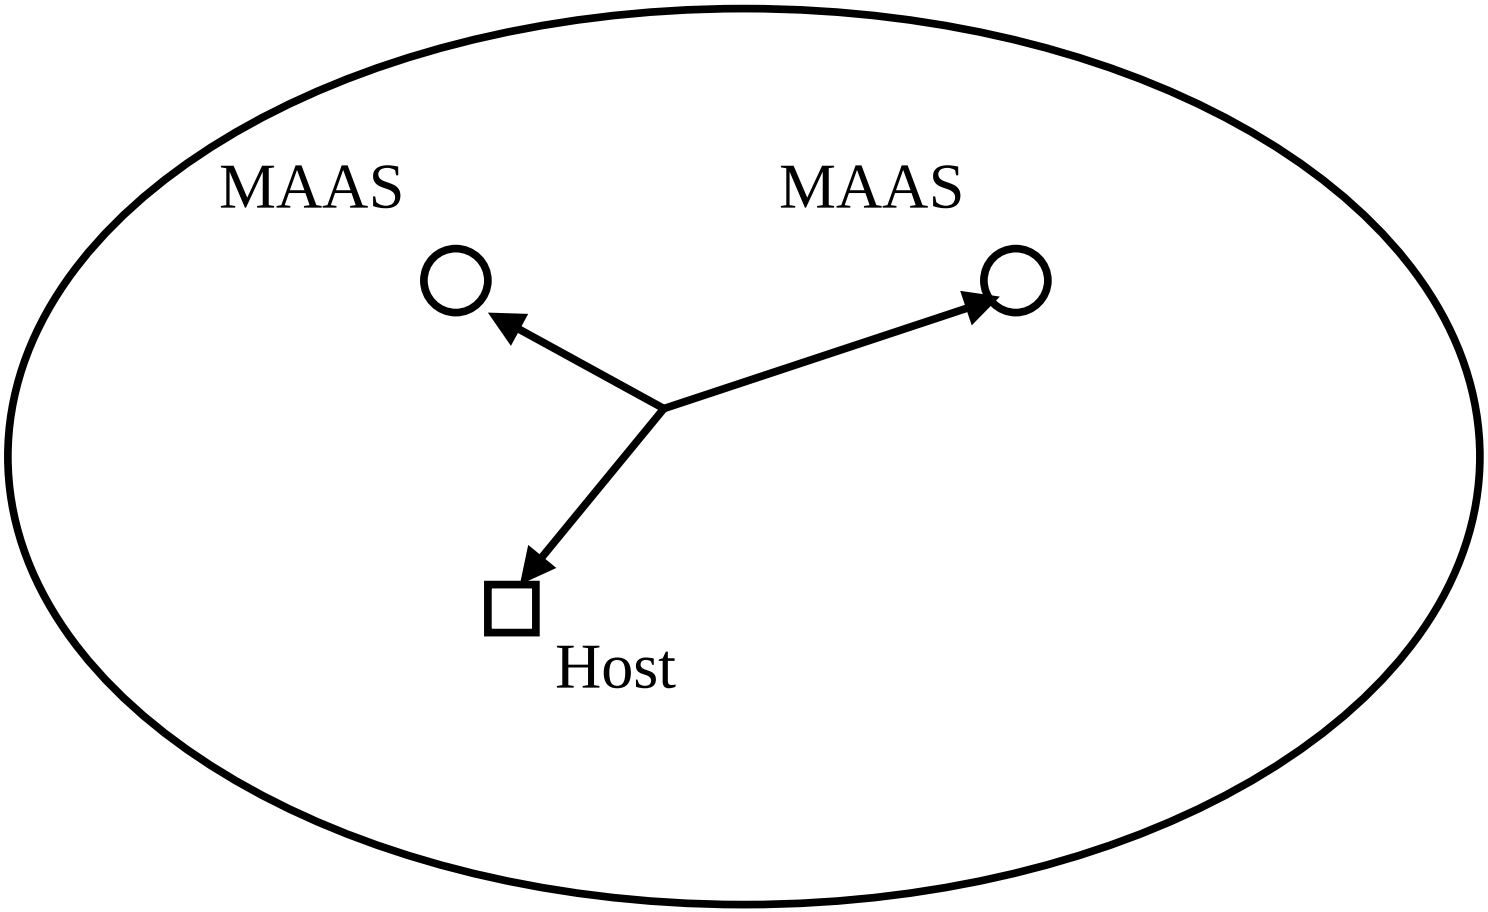

MDHCP Request-Response

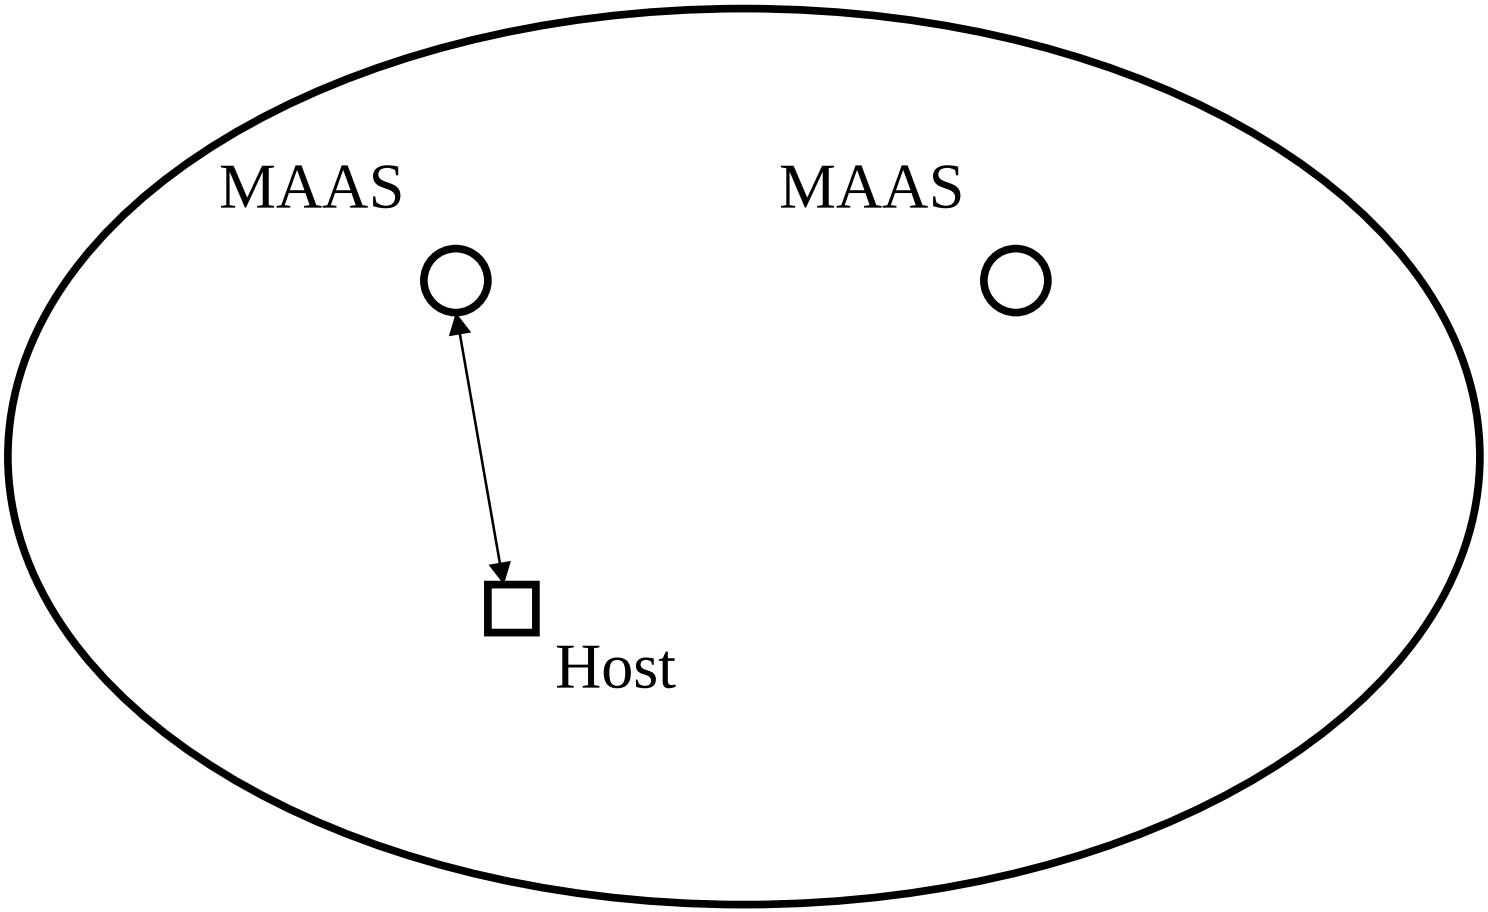

MDHCP

- Host-server
- Similar to DHCP
- Multicast Server Discovery
- Unicast Request-Response
- Fast and low-bandwidth
- Assumes coordinated servers

When to use malloc

- Use full mechanism for some set of global addresses; others may be allocated statically by IANA or in some other way
- For admin-scoped addresses, can use anything you like. We suggest using something like this. Do not need MASC for small scopes, though.

Document Status

- Architecture
  - Still needs minor revisions to reflect previous consensus
- MASC, AAP, MDHCP, and Abstract API
  - All recently updated
  - Resolved some open issues, some remain
  - MDHCP to WG Last Call 1/99
  - Other drafts to WG Last Call by 4/99?