Multicast Address-Set Claim (MASC) Update

MASC Internet Draft Status

- A pointer to the new draft just sent to the ML
- Several subsections added/changed
- A number of bug fixes
- One suggestion for change is still pending
- (Should be) ready to go for last call before end of July

Changes Since Last IETF

- Bootup Operation
- Leaf/Non-Leaf MASC Domain Operation
- Clock Skew Workaround
- Security Consideration
- "Sample Algorithms" Appendix added

Bootup Operation

- To create the appropriate sibling-parent association, establish connections to the PARENT nodes before the SIBLING nodes.
- To restore the internal state from more trustworthy source, establish connections to the PARENT nodes and/or INTERNAL PEER nodes before the connections to the CHILD and SIBLING nodes.

Leaf/Non-Leaf MASC Domain Operation

- A leaf MASC domain should advertise all of its managed addresses to the MAASs
- A non-leaf MASC domain should claim from itself (and compete against its MASC children domains) the space to advertise to the MAASs (default = 48 hours)
- The lack of claims by MASC child domains should trigger operation in "leaf mode", and vice versa

Clock Skew Workaround

Each claim's absolute timestamp is used to:

- Define collision winner.

  Collisions are presumably rare, hence the slow, "unfair" clocks are not a significant problem

- Estimate how long to keep a claim in a node's cache.

  Expiring sibling's PREFIX_IN_USE more than 48 hours earlier creates the potentials for clash. Keeping an expired claim for additional 48 hours will compensate for clock skew for up to 48 hours.

Security Consideration

- Trust your Parent and Internal Peers, but may drop internal state through Children and Siblings.
- Denial of Service attack (too many collisions) by a single node can be identified by all of its siblings, and ignore that node's claims.
- Denial of Service attack with multiple origin addresses can be prevented by accepting claims only through the parent, and "through" the claim node-originator itself.

"Sample Algorithms"

- Prefix allocation algorithms refined through (recent) simulations and briefly described in an Appendix
- Read the source code for algorithms details ;)

"Open" Issue

- Currently, siblings with more than one common parent can multiplex all UPDATEs over a single TCP connection
- Too complicated; negligible savings of few TCP connections
- Solution: open a new TCP connection between siblings for each common parent (simpler and easier to debug)

MASC Implementation Status

- Detailed testing, refining and bugs fixing of MASC processing code through simulations (*mascd* and *mascsim* share the same MASC-specific code).
- QUERY/RESPONSE debug messages added (describe them in APPENDIX?)
- MASC-AAP interface + AAP "client" in progress (it will be very helpful if someone already has MADCAP/AAP implementation)