

Multicast Lessons Learned

draft-mcbride-mboned-lessons-learned

Dino, Lenny, Mike

Background

- During a side meeting this summer it was mentioned that we really need a document discussing multicast lessons learned.
- The three of us met to discuss in Philly and came up with the draft as an initial effort with much more to go. This is not a BCP.
- Purpose is to help understand the previous development work to help understand the current work. Perhaps this will help with broader understanding of multicast and ongoing efforts for improvement.
- We start with DVMRP because that is really how multicast started on the Internet. We solved problems w/DVMRP and made design choices.
- IGMPv1 -> IGMPv3 and MLD
- DVMRP -> PIM-DM -> PIM-SM w/MSDP -> PIM-SSM

DVMRP

- The flooding and pruning of DVMRP was a good initial solution but we quickly realized that it wouldn't scale when using increasingly higher bit rates for multicast content.
- Using the network to discover sources was also something originally thought to be a good idea but later discovered to be resource and state intensive.
- DVMRP worked good for small scale deployments but began to suffer when deployed in larger multicast environments so we needed better solutions.

Shared and Source Trees

- Shared trees were designed to reduce state at a time when memory was scarce and expensive, while shortest path trees were simpler, and more optimal, but consumed more state.
- We had to fix the flood-and-prune problem that DVMRP had. We fixed that problem but didn't provide any explicit signaling from the source to discover them. So the multicast routing protocol discovered the sources (via the PIM shared-tree).
- Switching from one tree (shared) to the other (source) was a difficult routing distribution problem. Because as you joined the source-tree, you had to prune that source from the shared-tree.

All or Nothing Problem

- Every layer 3 hop between the sourcing and receiving end hosts must support a multicast routing protocol.
- On the global Internet this "all or nothing" requirement tends to create an insurmountable barrier.
- One approach is to develop solutions that facilitate incremental deployment and minimize/eliminate the need for coordination of multiple parties. Overlay networking is one such approach.
- For example, AMT [RFC7450] allows end users on unicast-only networks to receive multicast content by dynamically tunneling to devices (AMT Relays) on multicast-enabled networks

So many others..

- MSDP, Auto-RP, BGMP, MOSPF... but we are focusing on major issues/lessons learned.
- Should we work on this as a wg? pim or mboned?