draft-ietf-mboned-interdomain-peering-bcp

09 (prague) – 14 (singapore)

P. Tarapore, R. Sayko, G. Shepherd, T. Eckert, R. Krishnan

IETF'100 Singapore, November 2017

Toerless Eckert, Huawei (Futurewei Technologies USA)

tte+ietf@cs.fau.de

Changes

- 09 – 11:
  - Mentioned transit domains are out of scope, some textual refinements, mentioning BR with tunnels (GRE/AMT) can be anywhere in domain, reference updates.
- -11 passed to IESG for review
  - https://datatracker.ietf.org/doc/draft-ietf-mboned-interdomain-peering-bcp/ballot/
    - Results in a long list of DISCUS from Alissa Cooper, Ben Campbell, Spencer Dawkins, Mirja Kuehlewind, Kathleen Moriarty, comments (no discus, minor) from Adam Roach, Eric Rescorla
- -12 converted -11 into XML and github
  - No content changes, but lots of format changes
  - Terrible work 'txt2xml' + lots of hand editing (thanks Henrik Levkowetz)
  - Always start with XML if you can in future drafts…
  - XML: makes rfcdiff easier (important for reviewers to see diffs on their comments), hopefully also RFC editor queue processing faster

Changes

- 12 – 14: http://tools.ietf.org/tools/rfcdiff/rfcdiff.pyht?url1=https://raw.githubusercontent.com/toerless/peering-bcp/master/draft-ietf-mboned-interdomain-peering-bcp-12.txt&url2=https://tools.ietf.org/id/draft-ietf-mboned-interdomain-peering-bcp-14.txt

- Clarified scope & assumptions in intro, refer to p2p “private peering” as default case, PIM-SSM/BGP +AMT/GRE as the sete of protocols used across peering.
- Tunnel benefits: partial upgrade, incremental extensions, 3.4 – ability to introduce multicast without AD-2 support, …
- Detailing how AMT relay discovery/selection is still ongoing work
- Various other textual improvements through review questions...

ASCII graphics for every option now

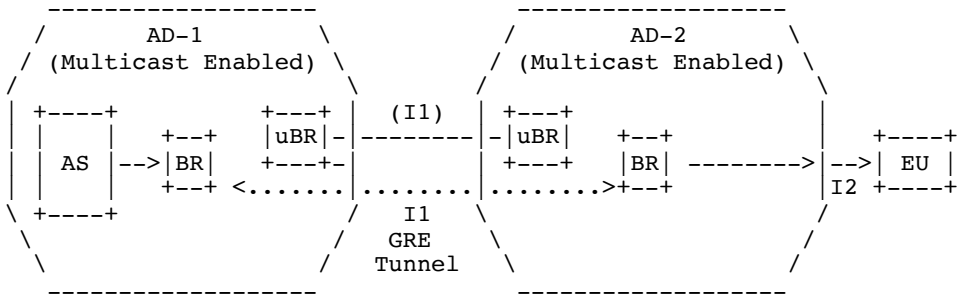

AD = Administrative Domain (Independent Autonomous System)
AS = Application (e.g., Content) Multicast Source
uBR = unicast Border Router - not necessarily multicast enabled
may be the same router as BR
BR = Border Router - for multicast
I1 = AD-1 and AD-2 Multicast Interconnection (e.g., MBGP)
I2 = AD-2 and EU Multicast Connection

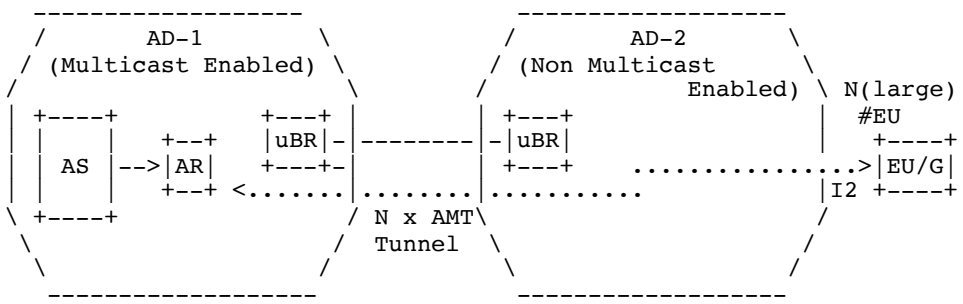

AS = Application Multicast Source
uBR = unicast Border Router - not multicast enabled,
otherwise AR = uBR (in AD-1).
AR = AMT Relay
EU/G = Gateway client embedded in EU device
I2 = AMT Tunnel Connecting EU/G to AR in AD-1 through Non-Multicast Enabled AD-2.

Changes

- 4.1.1 bandwidth management
  - TSV AD feedback response.
  - References to BCP41 and BCP145 (UDP multicast congestion control etc..)
  - Description of controlled vs. non-controlled network
  - Example risk with inelastic traffic when AD-1 assumes controlled network, but traffic is put into uncontrolled “best effort” in AD-2
  - Noting this is not an IP multicast issue but a problem of inelastic video app (same if VoD unicast would be inelastic)
  - Summarizing receiver rate adaption in multicast and risk of (S,G) changes create higher state maintenance performance requirements than traditional “inelastic IP multicast”.
  - Tunnels across third-party AS -> traffic MUST be rate adaptive unless third-party AS contracted for inelastic traffic.

Changes

- 4.1.1 public peering – multiple AS connecting via L2 LAN
  - Describe basic problem introduced by PIM assert
  - Describe solutions (tunnel across LAN, single upstream AS, federated / coordinated upstream AS).
- 4.3.2 / 4.3.3 / 4.6 inter-AS management interactions
  - Clarified workflow: sub -> AS1 -> AS2… why ? AS1 has content level relationship with sub, AS2 may not have any content level idea. Content relationship as keyword for no provider here is “spying more on sub data” than a unicast app provider would be able to do.
  - AS1 needs to collect more info about multicast from AS2 than in unicast because it can not infer loss in AS2 like it can in unicast (no loss feedback like in TCP)
  - Same concern for accounting (needs to happen after replication to sub)

Changes

- 5. troubleshooting
  - Mtrace / traceroute great … alas, they do not go through AMT tunnel (bummer) - may need more troubleshooting via other means.
  - [IMHO: If AMT becomes more widely deployed, bight want to think of mtrace over AMT spec extension]

- Security considerations
  - DoS attacks against state/bandwidth
    - Same as intradomain: limit amount of state/sub (good enough ?)
    - More difficult interdomain: Know which (S,G) are carrying valid traffic. May need to pass programming info from AS1 -> AS2 administratively to be put into ACL in AS2

Changes

- Security considerations
  - Content security
    - Not relevant for FTA content, but for content requiring DRM
    - Unicast: sender can prohibit receiver to get content (filter). Even with encrypted content, filtering is another key required security level – key can be cracked longer term or shared easily across network (like satellite TV keys are shared by hackers).
    - Intradomain solutions for filtering (S,G) per subscriber or sub-profile standard deployed in IPTV deployments. Interdomain solutions for this have not been well defined (not mentioned, but we tried to work on this 10 years ago in IETF).
  - Peering encryption
    - Prohibit leakage of content if risk of third party listening exists
  - Operational aspects
    - Info shared between AS1/AS2 may need to be protected: exposing (S,G) to content mappipng publically via databases may open additional attack vectors (eg.: DNS SSM mapping – not written into draft but not well secured exposure point).
      - Make operational data goes ONLY across peering point, encrypt peering point (“inband with actual media).
    - Existing text for token authentication, security breach mitigation plan

Changes

- Privacy considerations
  - AS1 has content relationship, so even if it was doing only unicast, it would know a lot about subscriber behavior. Multicast does not change this.
  - Only multicast novel privacy exposure: AS2 can likely deduce what content a subscriber watches if it wants (correlation based). In unicast it wouldn't be able to know this.
  - Is this bad or good ? Depends on content.
    - Most content, subscriber would like to see AS2 to provide good quality for even OOT content from AS1 and AS1 might even explicitly share info with AS2 (programming info).
    - Other type of content subscriber would not like this (e.g.: adult). Possible solution to bring multicast to same level of privacy as unicast: Make AMT tunnels encrypted, tunnel across AS2.

- THE END