IP Multicast
in Broadcasting World
-
Lessons Learned

9.6.2016 TREX Workshop
Antti Leimio
System Designer, Digita Oy
Twitter @AnttiLeimio

DIGITA

Background

- Broadcasting is one-to-many WAN distribution
  - Broadcast/multicast at its best
- Stream processing in Datacenter
  - Headend systems need multiple copies of streams
- Contribution signals and playout
  - More and more global delivery
- Traditional TV, radio is mostly static, always on service
  - Model is changing to more dynamic, customized
- Multicast routing old thing
  - PIM-SM 1998, PIM-SSM 2006
  - Digita: Mobile TV (DVB-H) 2006, MVPN since 2009, NG-MVPN since 2012
- IPv4 only, IPv6 feels impossible

DIGITA

Protocols and signaling

- Basic multicast routing: PIM-SM + IGMP
  - PIM = router to router
  - IGMP = host to router
  - Underlay: Unicast and RPF
- Use source-specific: PIM-SSM + IGMPv3
  - Receiver explicitly joins to group and source IP => Host does it all
  - Simple, easy, more secure and stable
  - No dynamic RP, register process, RPT/SPT switchover, MSDP
- SSM limitations
  - IGMPv3 host support? => Router SSM-mapping => Still PIM-SSM routing
  - Many-to-many, dynamic applications => Depends on application

DIGITA

Addressing

- Whole range 224/4 is usable with some restrictions
- Life is easier with default ranges:
  - 239/8 ASM
  - 232/8 SSM
- Plenty of private addresses to spend
- Groups 0, 255 could be issue somewhere
- Overlapping macs
  - Plan first two octets, use converter tools
- Use IPAM and DNS names

Multicast IP 239.1.1.1 converts to:
MAC address 01:00:5e:01:01:01

Matched multicast IP group addresses

224.1.1.1
224.129.1.1
225.1.1.1
225.129.1.1
226.1.1.1
226.129.1.1
…
232.1.1.1
232.129.1.1
233.1.1.1
233.129.1.1
234.1.1.1
234.129.1.1
235.1.1.1
235.129.1.1
236.1.1.1
236.129.1.1
237.1.1.1
237.129.1.1
238.1.1.1
238.129.1.1
239.1.1.1
239.129.1.1

DIGITA

Addressing

- ”Public” space is problem
  - Needed for inter-domain streams
  - Very limited space
  - 233/8 GLOP (16-bit AS) = 256 addresses per AS
  - 234/8 RFC6034 Unicast prefix-based (32-bit AS) => not usable
- SSM can reuse same group
  - Same group with different source
  - Complexity in troubleshooting
  - Possible routing/switching/host issues

Routers

- Redundancy in LAN
  - Forwarder = PIM Assert winner (Lowest AD, metric, then highest interface IP wins)
  - IGMP Querier = Lowest IP address queries hosts and keeps table
  - VRRP/HSRP not involved
  - Could be different than unicast topology
- IGMP Querier in LAN
  - L2 LAN needs querier on a switch (IGMP capable device)
  - Configurable (but global option)
  - Assign lowest IP to router interface
- IGMP is the complex part!
  - Hosts implementation/configuration issues => Most problems
  - Complex set of timers => Don’t touch

Switches

- IGMP snooping is your friend
  - Turn it on
  - Reduced flooding
  - Visibility to L2 domain: Groups and ports
- Problems:
  - Lack of IGMPv3 source specific support
  - Snooping support (per vlan): Blades, virtualization, low end switches
  - Flooding on uplink ports if vlan configured at both ends
- Storm-control
  - Multicast traffic must be excluded

MPLS

- Native MPLS forwarding: LSP, L2VPN, L3VPN
  - mLDP: MP2MP, better scalability, one protocol
  - RSVP: P2MP, TE, FRR, scalability issues
- RSVP-TE P2MP LSP: Multicast tree of your choice
  - Dynamic, static or mix
  - Active-active, active-standby
  - Resource reservation, QoS
  - Link-protection
- L2VPN
  - H-VPLS
  - Control and visibility at transport layer?

DIGITA

MPLS L3 MVPN

- Legacy: Draft-Rosen PIM-tunnels
- Use robust NG-MVPN
  - BGP signaled PIM free core
  - VPN specific PIM islands on PE
  - RSVP-TE, FRR options
  - VPN Extranet support
- Issues
  - I-PMSI flooding => Use S-PMSI => More RSVP tunnels
  - RSVP scalability => RSVP Refresh Reduction, disable unnecessary provider-tunnels
  - Extranet is complex

Inter-domain, static

- Static push model, always on
  - PIM-DM could do that, support?
  - Independent domains => No shared protocol state, routes
  - Static igmp join (router + switch)
  - Active-active model, same or different groups on redundant connections
  - Stream selection on receiving device
  - Content rights => Filtering outgoing streams => Dynamic and labour intensive
  - Router HW with ACLs, Firewalls?
  - MPLS MVPN extranet use case: Distribution hub VPN -> customer VPNs
  - AS-based ”public” GLOP addressing

DIGITA

Inter-domain, dynamic

- Dynamic PIM routing
  - “PIM peering”, shared state
  - Receiver subscribes streams
  - Dynamic routing => Possible instability
  - Network level redundancy with PIM
  - Straightforward configuration with PIM-SSM
  - Need content filtering also

DIGITA

OAM

- Multicast is something mysterious
  - Fundamentally different than unicast, like overlay
  - Knowledge is needed
  - Basic operation and configuration is simple and easy
- Visibility and monitoring
  - Stream monitoring is easy, just pick it
  - Use L3 segmentation to see what is happening
  - Router’s PIM, IGMP, MVPN has good tools
  - Data plane measurement: Video monitoring features
  - Lack of big picture
- Operating
  - Not much to do, PIM/IGMP just works
  - Biggest problems with IGMP hosts

```
JUNIPER> show multicast route instance VRFX group
233.188.171.197 detail

Group: 233.188.171.197
   Source: x.x.x.x/32
   Upstream interface: ae1.251
   Downstream interface list:
      ae1.1027 ae1.1017 ae1.1016 ae1.1015 ae1.1014
ae1.1527 ae1.1526 ae1.1522 ae1.1252
   Session description: GLOP Block
   Statistics: 1417 kBps, 1045 pps, 6247493158 packets
   Next-hop ID: 1048608
   Upstream protocol: MVPN
```

```
JUNIPER> monitor traffic interface ae1.151 matching igmp
09:27:19.518347  In IP 10.10.10.2 > 232.3.1.36: igmp query v3 [max resp time is 10
units or 1.0s] [gaddr 232.3.1.36, 1 source(s)]
09:27:19.620059  In IP 10.10.10.37 > 224.0.0.22: igmp v3 report, 2 group record(s)
09:27:19.682034  In IP 10.10.10.97 > 224.0.0.22: igmp v3 report, 1 group record(s)
09:27:19.816485  In IP 10.10.10.123 > 224.0.0.22: igmp v3 report, 1 group record(s)
09:27:19.960297  In IP 10.10.10.95 > 224.0.0.22: igmp v3 report, 2 group record(s)
```

DIGITA

Future

- Multicast is love/hate thing
- Use cases: Efficient realtime UDP transport
  - Live streaming: Broadcasting, IPTV, CDN origin -> edge
  - Mass updates
  - Mostly inside one administrative domain
- Show stoppers
  - Internet => Unicast world
  - Personalization/customization of streams => More like unicast
  - HW/SW limitations, merchant silicon => Only big boxes support MPLS features
  - Virtualization => Limited MC support
  - Public address space => Needs IANA level actions
  - Inter-operability

DIGITA

Future

- New possibilities
  - SP Network: AMT, BIER
  - Mobile: LTE Broadcast, 5G
  - SDN: Openflow, BGP-LS, etc.
- Internet will never have multicast
  - Multicast remains backend/backhaul transport method
  - Still need for multicast peering? IXPs?
- Tools must also evolve
  - Chicken-egg-problem
  - Maybe SDN and analytics?

DIGITA

Thank you!

DIGITA