IPv6 multicast

Stig Venaas
venaas@uninett.no

UNINETT

Contents

- State of multicast
- What is new with IPv6 and multicast?
- IPv6 multicast addresses and scopes
- IPv6 multicast protocols
- Interdomain IPv6 multicast
- Current deployment
- IPv6 multicast applications
- Conclusions

UNINETT

What is IP multicast?

- Usually an IP packet is sent to one specific host
  - The IP destination address specifies which host
- With IP multicast, an IP packet is sent to a group of hosts
  - The IP destination address is a group and not a host address
  - IPv4 multicast addresses, class D. 224.0.0.0 – 239.255.255.255
  - The group can contain any number of hosts (0 to infinity)
  - The group members can be anywhere
- Multicast packets will be replicated by routers where needed
  - Routers keep track of which interfaces should forward the packet
  - The same multicast packet is ***never sent twice on the same link***, hence the bandwidth used on a specific link is ***independent of the number of receivers***

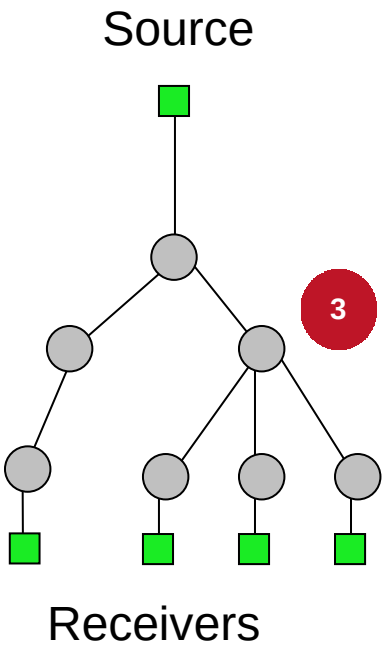

Multicast – state of the union

- Current state of multicast is not as good as some of us had hoped
- Not so much multicast on the Internet but in several closed networks
  - IPTV with multicast is becoming a big success
- Source-Specific Multicast (SSM) offers simplified and in some ways better multicast service, not much used yet
- Not that many applications because multicast is not available in general
- IETF is working in Automatic Multicast Tunneling
- P2P applications often do some kind of multicast on the application layer
- IRTF SAM (Scalable Adaptive Multicast) wg looking at generic solution where native multicast is used when available, and application (p2p) multicast when not

UNINETT

What is new with IPv6 and multicast?

- All IPv6 hosts and routers support link-local multicast
  - Multicast used instead of broadcast
  - Also support MLD (like IGMP for IPv4)
- Better defined scoping to control who receives what
- The size of the addresses is used to allow better ways to assign addresses, less risk of address conflicts
- We will see that it scales better for interdomain use
- Most routers can do multicast with no multicast related configuration

UNINETT

IPv6 multicast addresses

11111111	Flags	Scope	Group ID
8 bits	4 bits	4 bits	112 bits

- There are several types of multicast addresses
- Unicast prefix based
  - Including one embedding the RP addresses
  - Makes it easier to create unique multicast addresses
- SSM addresses

UNINETT

IPv6 multicast protocols

- MLD is used for hosts to report to routers what they want to receive
  - Mostly similar to IGMP for IPv4
- MLDv1 supports joining groups (ASM)
  - Like IGMPv2
- MLDv2 supports joining groups and sources (SSM)
  - Like IGMPv3
- All IPv6 systems (should) support MLDv1
- MLDv2 supported by Linux and Windows Vista
- PIM-SM is the most widely used routing protocol
  - Just like it is for IPv4
- Multiprotocol BGP used for exchanging routes for RPF
  - Like for IPv4
- One big difference is that there is no MSDP, so as we shall see, interdomain multicast is different

UNINETT

PIM-SM and Rendezvous Points

- Interdomain multicast routing is usually done with a protocol called PIM-SM
- PIM-SM requires an RP for source discovery
- All routers must use the same RP and somehow know its address
- Initially packets from a source will be sent to the RP
- When a host joins a group, join messages are sent hop by hop towards the RP
- The RP serves as a meeting place between sources and receivers
- This works well within a site or a single administrative domain
- But we don’t want one single central common RP for all multicast on the Internet
- So we want something that can work across administrative domains

UNINETT

Interdomain IPv4 multicast 1/2

- Different administrative domains each have their own RP(s), even for global groups
- Sources and receivers in a domain will just send to or join the domain's RP
- So how can we have communication between domains?

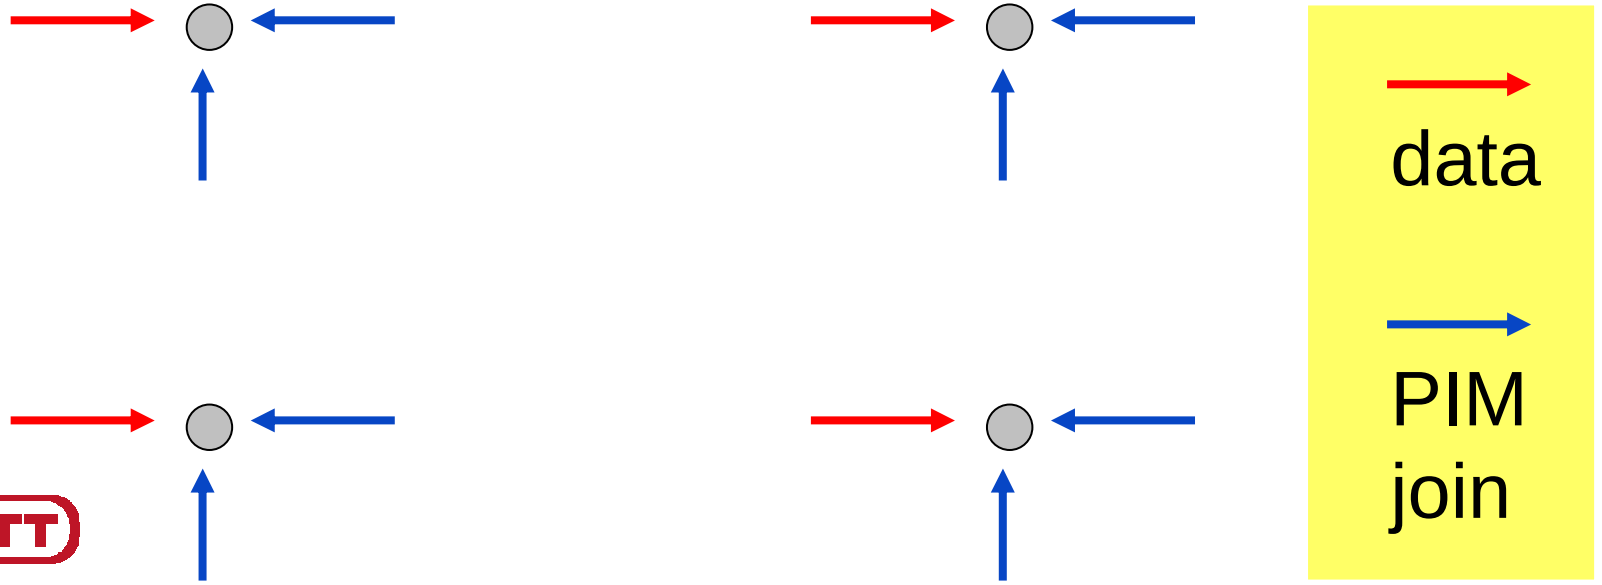

UNINETT

Interdomain IPv4 multicast 2/2

- MSDP (RFC 3618) is used to solve this
- Network of MSDP peerings between the domains RP's
- When a source starts sending in one domain, MSDP will send source announcements to all other RPs. Repeated periodically
- When someone joins in some domain, a source specific tree is built from that domain's RP to the source in the remote domain

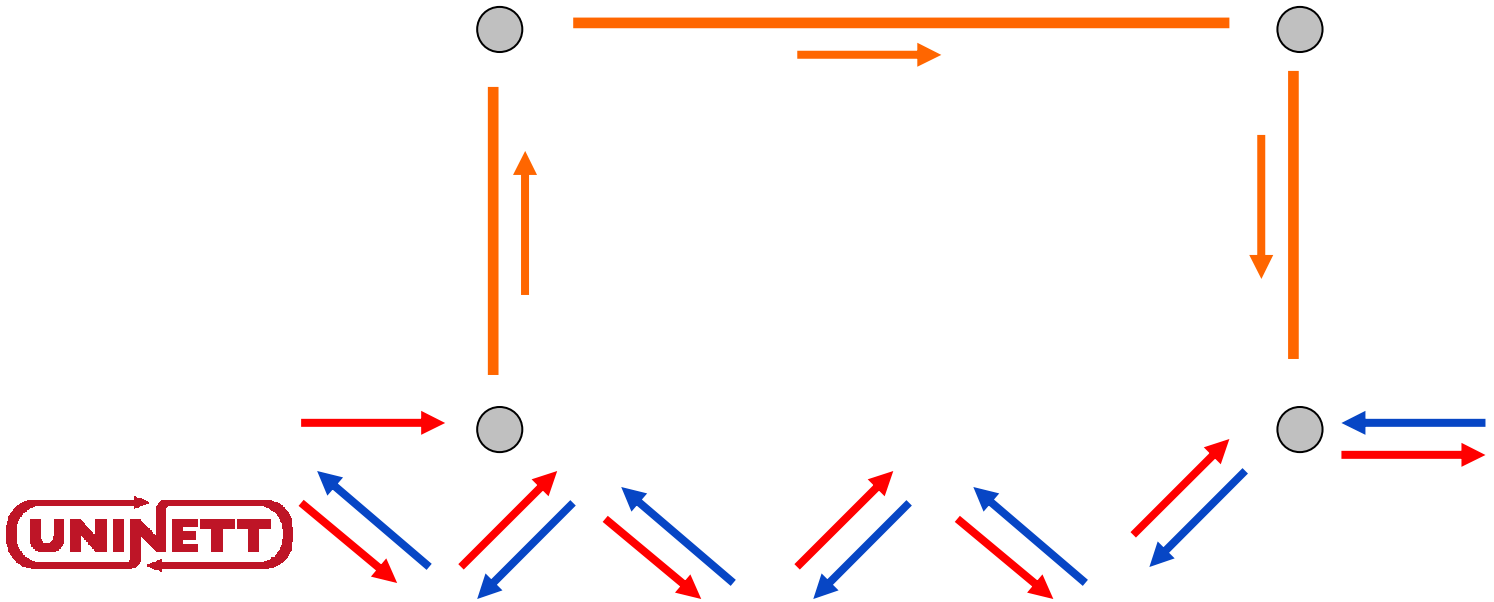

IPv6, MSDP and Embedded-RP

- MSDP doesn't scale with large scale deployment since source information is flooded to all the RPs
- Hence the IETF decided not to have MSDP for IPv6
- For IPv6 there is something called Embedded-RP (RFC 3956)
- It defines a specific way to create multicast group addresses where the RP address is encoded into the group address
  - An embedded-RP address starts with ff70::/12
  - Flags value of **7** means embedded-RP
  - E.g. ff7e:140:2001:700:f000:100:1234:beac has the RP 2001:700:f000:100::1
- Only a new way to map from group to RP. The main point is that it allows for a large number of RPs that can be practically anywhere in the Internet. They do not need to be preconfigured in the routers, routers automatically use the right RP
- Someone hosting or initiating a multicast session can pick a group address with their RP address encoded inside
- Everyone joining or sending to their session will then use their RP

Interdomain IPv6 multicast

- We have four domains, each with their own RP used for sessions they are hosting or initiating
- When someone joins in a domain, a shared tree is built from the last-hop router (where the joining host is) towards the RP of the host/initiator
  - The RP address is derived from the group address

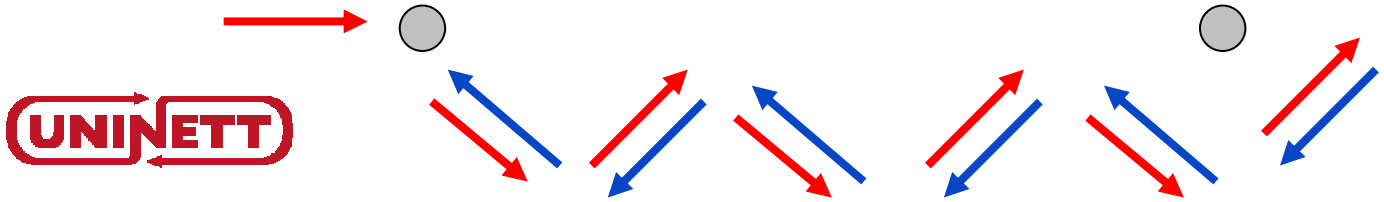

IPv4 – IPv6 comparison

- The main philosophy behind MSDP is to avoid one single common RP in the Internet, and to avoid relying on a 3rd party's RP
- With Embedded-RP we solve both of these
  - Provided one of the parties in a session picks a group address specifying their own RP
- A technical difference is that with MSDP (and also SSM), there is only (S,G)-joins between domains
- With embedded-RP, an RP is shared by multiple domains, so there will also be (*,G)-joins between domains

UNINETT

Embedded-RP issues

- All routers on the paths between the RP and sources/receivers must support it
- A user in one domain is now using an RP in a remote domain instead of a local one. This may make it harder to debug multicast problems, since the user and the RP are in different domains
- No changes required in hosts or applications, except:
  - How does an application or a user know which group address to use? They should not need to care about RPs. Admins need to have a way to tell users or applications which group address to use or which range to pick addresses from
- How to control usage of RPs?
  - Is it a problem if other people pick group addresses with your RP address encoded?
  - Is there a reason for someone to do this? Does it matter, can it be prevented?
- What about sessions that have no owner?
  - E.g. session discovery SAP/sdr. Don't want to rely on one particular RP for this
- These problems are being worked on

What about SSM?

- SSM (Source-Specific Multicast) available for both IPv4 and IPv6
  - There are no differences here
- Some believe only SSM is needed for interdomain multicast
- SSM simplifies multicast signaling in the network
  - No need for RPs, PIM register, switching between shared and source-specific trees…
- But very difficult for multi-party applications
  - E.g. conferencing where everyone is a source and everyone needs to know the IP addresses of the others
- SSM is supported by very few applications and systems
  - Edge routers and hosts need to support MLDv2
  - Hosts need to support RFC 3678, which is the API for specifying source filters
  - Applications need to be changed to support this API

Deployment of IPv6 multicast

- IPv6 multicast has been deployed natively in Abilene, GEANT and NORDUnet, as well as in several NRNs (incl. FUNET and UNINETT)
  - These also have native IPv6 multicast peerings
  - GEANT has peerings with about 10 NRNs
  - Abilene has peerings with GEANT, CA*net 4, DREN and AARNet
- There is also a world-wide network called M6Bone. This was started by RENATER 6 years ago
  - Anyone can join this to experiment with IPv6 multicast, see http://www.m6bone.net/

UNINETT

M6Bone – Europe

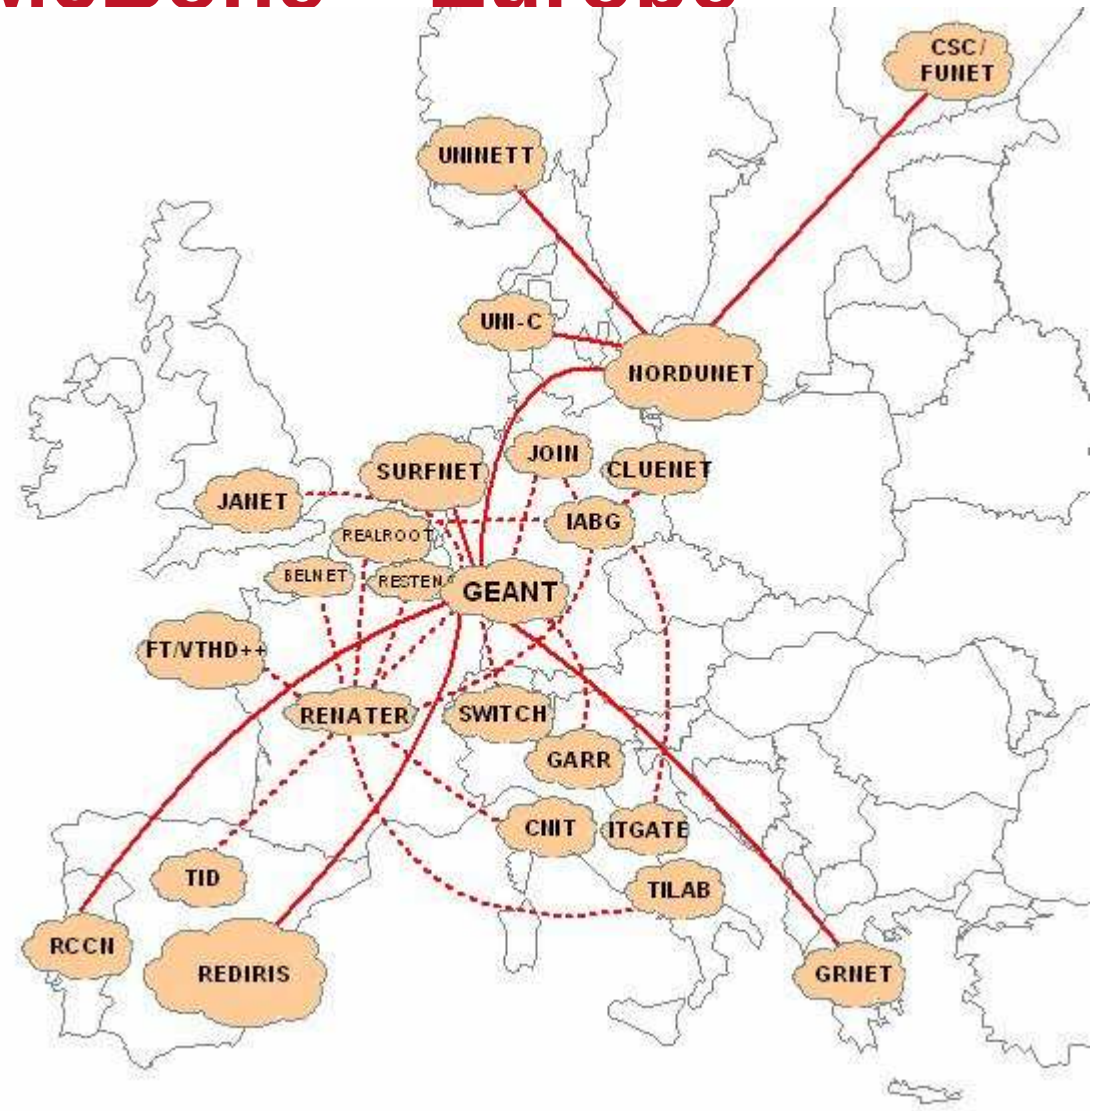

M6Bone – The World

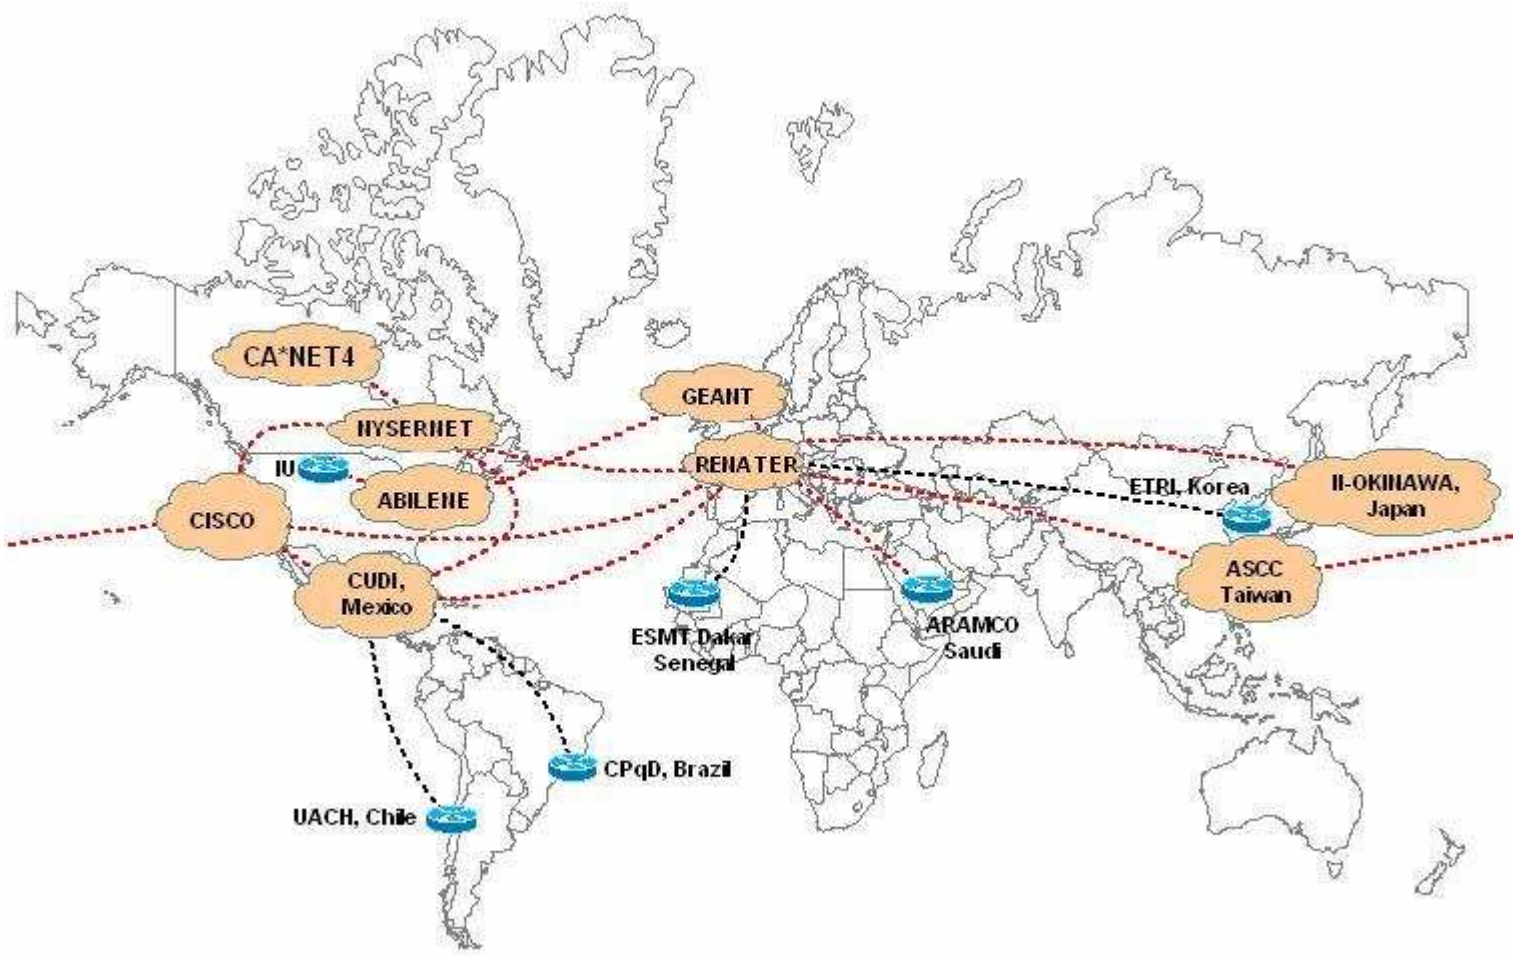

Deployment of IPv6 multicast

- IPv6 multicast with a single central RP has been working in the M6Bone for 6 years. Still in use, but as we've discussed, it's not a scalable solution
- Embedded-RP and SSM are being actively used in NORDUnet, GEANT, Abilene and several NRNs (and end sites)

IPv6 multicast applications

- Mbone tools – vic/rat etc
  - IPv6 multicast conferencing applications
  - http://www-mice.cs.ucl.ac.uk/multimedia/software/
- VideoLAN (vlc) – http://www.videolan.org/
  - Video streaming, also IPv6 multicast. Server and client
- DVTS – http://www.sfc.wide.ad.jp/DVTS/
  - Streaming DV over RTP over IPv4/IPv6
- Mad flute – http://www.atm.tut.fi/mad/
  - Streaming of files using multicast (IPv4/IPv6 ASM/SSM)
- Monitoring/debugging tools
  - Dbeacon, http://dbeacon.innerghost.net/
    - An alternative to NLANR beacon
  - ssmping/asmping – http://www.venaas.no/multicast/ssmping/
    - ping-like tool for testing IPv6/IPv4 ssm/asm connectivity

UNINETT

Conclusions

- IPv6 multicast is being deployed in current production networks
  - At least academic networks
- IPv6 multicast is easier to deploy and more scalable than IPv4
  - All IPv6 hosts and routers need to support at least link local multicast
  - Well defined scoping architecture
  - Embedded-RP scales better than MSDP
  - No configuration necessary for non-RP routers
- We expect to see SSM and Embedded-RP being the IPv6 solutions for multicast across multiple domains in the Internet
- Some argue we only need SSM
  - This makes it very simple for the network for both IPv4 and IPv6, but adds complexity to applications
- There are several applications and tools available, but need more