

# IP Multicast

“The Good, the Bad and the Ugly”

*RIPE 52 - Istanbul*

*John Lyons  
Zen Internet, UK*



# Background

Spent the last few years designing, implementing & troubleshooting multicast networks in :

- Research and education networks
- “Multi play” ISP delivering IPTV over FTTH/ETTH/ADSL2+
- Traditional UK ISP with native multicast delivery for BBC trial

**Disclaimer : All views/opinions/rants are my own personal views/opinions/rants**



# Multicast deployment

- Multicast deployment has historically been very poor
- Lots of reasons : “the code’s not stable”, “there’s no business case”, “no-one else is deploying it”, “I don’t have time to learn it” ...
- In the last 12 - 24 months this has **finally** begun to change



# Who's deploying it ?

- “Multi Play” ISP's
- Research and education institutions
- Traditional ISP's
- Large scale content providers
- Financials
- 3G Wireless providers



# I want to deploy it !

- Great ! :-)
- I could paint a rosy picture but I'd be lying
- Multicast has it's fair share of problems
- However, with good preparation & design many of these problems can be worked around
- My experiences with multicast have fallen into 3 categories - "The Good, the Bad and the Ugly"



# The Good ...

- General benefits
- Well designed ASM networks
- SSM
- PIM Bidir
- Configuration
- Results




# The Good ...

- General benefits
  - Well designed ASM networks
  - SSM
  - PIM Bidir
  - Configuration
  - Results




# General Benefits

Unicast



Multicast





# The Good ...

- General benefits
- Well designed ASM networks
- SSM
- PIM Bidir
- Configuration
- Results



# Well designed ASM

- Any Source Multicast (ASM) is the most widely deployed form of multicast today (Using PIM-SM, some legacy PIM-DM)
- It works best on well designed, hierarchical networks
- Relies on RPF and PIM-SM to make intelligent forwarding decisions
- Is now being used widely to deploy “closed garden” IPTV solutions



# The Good ...

- General benefits
- Well designed ASM networks
- **SSM**
- PIM Bidir
- Configuration
- Results



# SSM

- SSM is the answer to “one to many” applications
- No need for complications such as RP’s or MSDP
- No shared tree
- No address allocation issues
- Improved security
- Requires host kernel support for IGMPv3
- Still not widely deployed, mainly due to vendor support



# The Good ...

- General benefits
- Well designed ASM networks
- SSM
- PIM Bidir
- Configuration
- Results



# PIM Bidir

- The answer to “many to many” applications such as video conferencing & financial applications
- Relies on shared tree to forward all traffic
- Greatly reduces the amount of state which a router must store
- Is not widely deployed, yet
- Still relies on an RP (although only a logical RP) but not MSDP



# The Good ...

- General benefits
- Well designed ASM networks
- SSM
- PIM Bidir
- Configuration
- Results



# Configuration

- It's simple to configure :

## *Sample configuration for Cisco router :*

```
ip multicast-routing  
ip pim rp-address 10.193.0.1  
int x  
ip pim sparse-mode
```

## *MSDP gets more complicated but still relatively simple :*

```
ip msdp peer 85.91.0.6 connect-source Loopback2  
ip msdp cache-sa-state  
ip msdp originator-id Loopback2  
ip msdp mesh-group mesh-group-magnet-internal1 85.91.0.6
```



# The Good ...

- General benefits
- Well designed ASM networks
- SSM
- PIM Bidir
- Configuration
- Results



# Results

- A lot of networking projects give subtle results which are difficult to explain to management
- Deploying multicast usually results in streaming a DVD or TV channel network wide
- Never underestimate the wow factor :-)





# The good ...

- That all looks great
- Some reading and preparation and I'm ready !
- Well ... yes and no



# The bad ...

- ASM limitations
- Inter-domain ASM
- Troubleshooting
- Ongoing support



# The bad ...

- **ASM limitations**
  - Inter-domain ASM
  - Troubleshooting
  - Ongoing support



# ASM limitations

- Unfortunately basic ASM (PIM-SM, RP, MSDP, IGMP) has a lot of limitations
- Security (more in the “ugly” section)
- Address allocation for inter-domain multicast
- Lots of state
- SSM and PIM Bidir were built to address these problems



# The bad ...

- ASM limitations
- Inter-domain ASM
- Troubleshooting
- Ongoing support



# Inter-domain ASM

- Global ASM - MSDP used to maintain list of sources, MBGP used to maintain RPF
- Requires a large amount of state on RP and has potential to cause significant problems
- Inter-domain multicast is frustrating to troubleshoot and requires a lot of cooperation with other AS'
- Inter-domain multicast is not being tested heavily in production, there are very few global sources outside of the NREN's
- Consensus has been reached within operators that global ASM does not scale and is not the way forward



# The bad ...

- ASM limitations
- Inter-domain ASM
- Troubleshooting
- Ongoing support

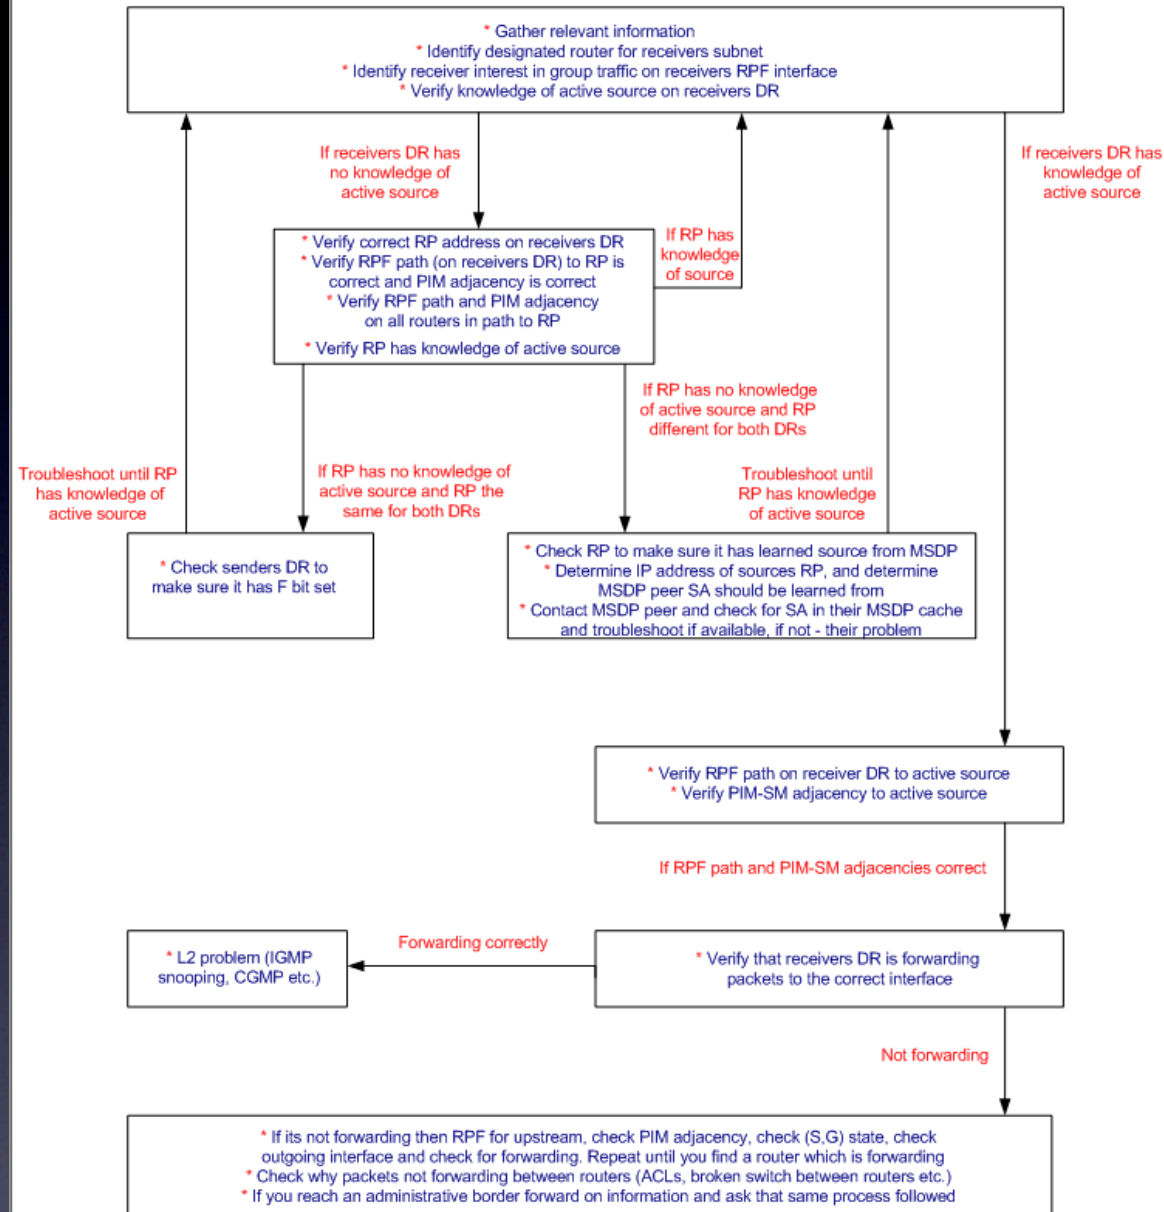


# Troubleshooting

- This is a separate talk all on its own
- Poses significant problems as it's receiver driven, is unidirectional in nature, has complicated rules and needs state debugging on all participating network devices
- It is possible to troubleshoot effectively but requires a significant initial effort to develop a step by step process
- Even still, in inter-domain multicast troubleshooting you will have to rely heavily on other people/networks



**Multicast troubleshooting – John Lyons  
(12/04/06) V2.1**



**Commands (Cisco) :**

```

Show ip route receivers interface
Show ip pim neighbor RPF interface
Show ip igmp groups RPF interface
Show ip mroute group-address source-address
Show ip pim rp mapping
  
```

**Commands (Juniper) :**

```

Show pim join extensive
Show multicast rpf source/RP
Show igmp group group address
Show igmp interface RPF interface
Show multicast usage
  
```



# The bad ...

- ASM limitations
- Inter-domain ASM
- Troubleshooting
- Ongoing support



# Ongoing Support

- Finding (and keeping) multicast literate engineers is difficult
- If you rely on multicast to deliver a critical service you will require experienced multicast engineers
- Training current staff can be difficult, but it is possible with the right methods



# The Bad ...

- That still looks fine
- I'm ready, give me the books and enable/config access !
- Well, there are some other things you should know ...



# The Ugly ...

- MAC address 32:1 overlapping
- Security
- Software bugs
- Vendor support
- Tunneling
- Inappropriate usage



# The Ugly ...

- MAC address 32:1 overlapping
  - Security
  - Software bugs
  - Vendor support
  - Tunneling
  - Inappropriate usage



# MAC Addresses

- Multicast L3 addresses are 28 bits long (/4), Special MAC addresses reserved for multicast, but due to lack of funding for OUI's, only 23 bits long
- There are 32 IP multicast L3 addresses with the same MAC address
- Host takes care of MAC address crossover but has to de-capsulate packet
- Particularly ugly if you choose to use multicast addresses which overlap with special case addresses such as 224.0.0.0/24 (224.0.0.1 or 239.0.0.1 = 01:00:5e:00:00:01)
- Renumbering is never fun



# The Ugly ...

- MAC address 32:1 overlapping
- Security
- Software bugs
- Vendor support
- Tunneling
- Inappropriate usage



# Security (Or lack thereof)

- Lots of security concerns (Mainly DoS through excess state)
- Bogus PIM-SM DR/IGMP querier elected
- Host joining too many groups
- Host sending to many groups
- Host sending to existing group
- Most of the fixes involve ACL's or using rate-limiting which can inadvertently affect legitimate users



# The Ugly ...

- MAC address 32:1 overlapping
- Security
- **Software bugs**
- Vendor support
- Tunneling
- Inappropriate usage



Bug ID #CSCdx82485

Symptoms: Under rare circumstances, a router that is configured with Protocol-Independent Multicast (PIM) may pause indefinitely.

Workaround: Use a different Ethernet card, or **avoid using PIM**



# Software bugs

- Here's where it can all go wrong
- Multicast software bugs are numerous
- One particular problem in a University led me to the bug tool for a Cisco 6509 with an older software revision, it had 350 multicast bugs
- Choose your software versions and hardware vendors wisely
- For critical services, find a stable image and stick to it, don't move without rigorous multicast lab testing



# The Ugly ...

- MAC address 32:1 overlapping
- Security
- Software bugs
- **Vendor support**
- Tunneling
- Inappropriate usage



# Vendor support

- Many vendors visibly wince when asked about their multicast support
- Workarounds for PIM problems will include fixes such as turning off PIM to resolve the problem
- Vendors can take significant time to find reasons for faults and even longer to resolve issues
- Long term - Business requirements and reliance on multicast for critical revenue generating services will solve this problem
- Short term - Ask questions and choose wisely



# The Ugly ...

- MAC address 32:1 overlapping
- Security
- Software bugs
- Vendor support
- **Tunneling**
- Inappropriate usage



# Tunneling

- If you have non multicast aware boxes (particularly firewalls) you may need to tunnel around them
- Static mroutes are non-intuitive and ongoing maintenance of them is time consuming
- Need to ensure tunnel endpoints are capable of pushing high quality videostreams through tunnel interfaces
- Increases complexity for troubleshooting



# The Ugly ...

- MAC address 32:1 overlapping
- Security
- Software bugs
- Vendor support
- Tunneling
- Inappropriate usage



# Inappropriate usage

- Some people will try and persuade you that it's a good idea to run video conferences with many participants in many Universities around the world using interdomain multicast
- These people obviously **do not** troubleshoot multicast problems
- A survey of GRID users using AccessGRID showed that > 20% of users "usually cannot" or are "never able" to use multicast  
<http://www.andrewpatrick.ca/multicast-survey/index.html>
- Go figure



# Conclusions

- Multicast finally has some problems to solve
- Multicast adoption will continue to rise, particularly in “walled garden” environments
- Global multicast will rely on SSM
- With preparation and hard work you can use it to run production services
- We all need to work with vendors to ensure software and hardware problems are resolved and new protocols are supported



# Resources

- Developing IP Multicast Networks - Beau Williamson
- Inter-domain IP Multicast - Edwards, Giuliano, Wright
- “A methodology for troubleshooting IP Multicast”  
- Litvanyi / Nickless (NANOG 27)
- “In the Trenches with IP Multicast” - Greg Bell (LBLnet)
- Internet 2 Multicast workshops
- Vendor documentation



Thank you for listening

Q & A

[john@techvictim.net](mailto:john@techvictim.net)

[www.techvictim.net](http://www.techvictim.net)