

IP Multicast

“The Good, the Bad and the Ugly”

John Lyons
Magnet Networks
Network Development

Background

- 3 years designing, implementing & troubleshooting multicast networks
- Extensive experience with research and education multicast networks
- Experience in critical production environment with Magnet Networks in Ireland : Streaming digital TV over FTTH/ETTH and ADSL2+ using multicast
- **Disclaimer : All views/opinions/rants are my 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” ...
- All of these reasons were viable, however this is now /finally/ beginning to change

Who's deploying it ?

- Companies delivering IPTV
- Universities
- Financial companies
- Large scale content providers

I want to deploy it!

- Great!
- I could paint a rosy picture but i'd be lying
- There are problems, but with good preparation & design they can be worked around
- “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
- On badly designed, badly maintained networks it can fail spectacularly and in unexpected ways

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 Bi-dir

- 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 but not MSDP

The Good ...

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

Configuration

- It's trivial 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 IP Multicast
- Troubleshooting
- Ongoing support

The bad ...

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

ASM limitations

- Unfortunately basic ASM (PIM-SM, RP, MSDP, IGMP) has a lot of limitations
- Security
- Address allocation for inter-domain multicast
- State tables
- SSM and PIM Bidir were built to address these problems

The bad ...

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

Inter-domain Multicast

- 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 source outside on the NRENs

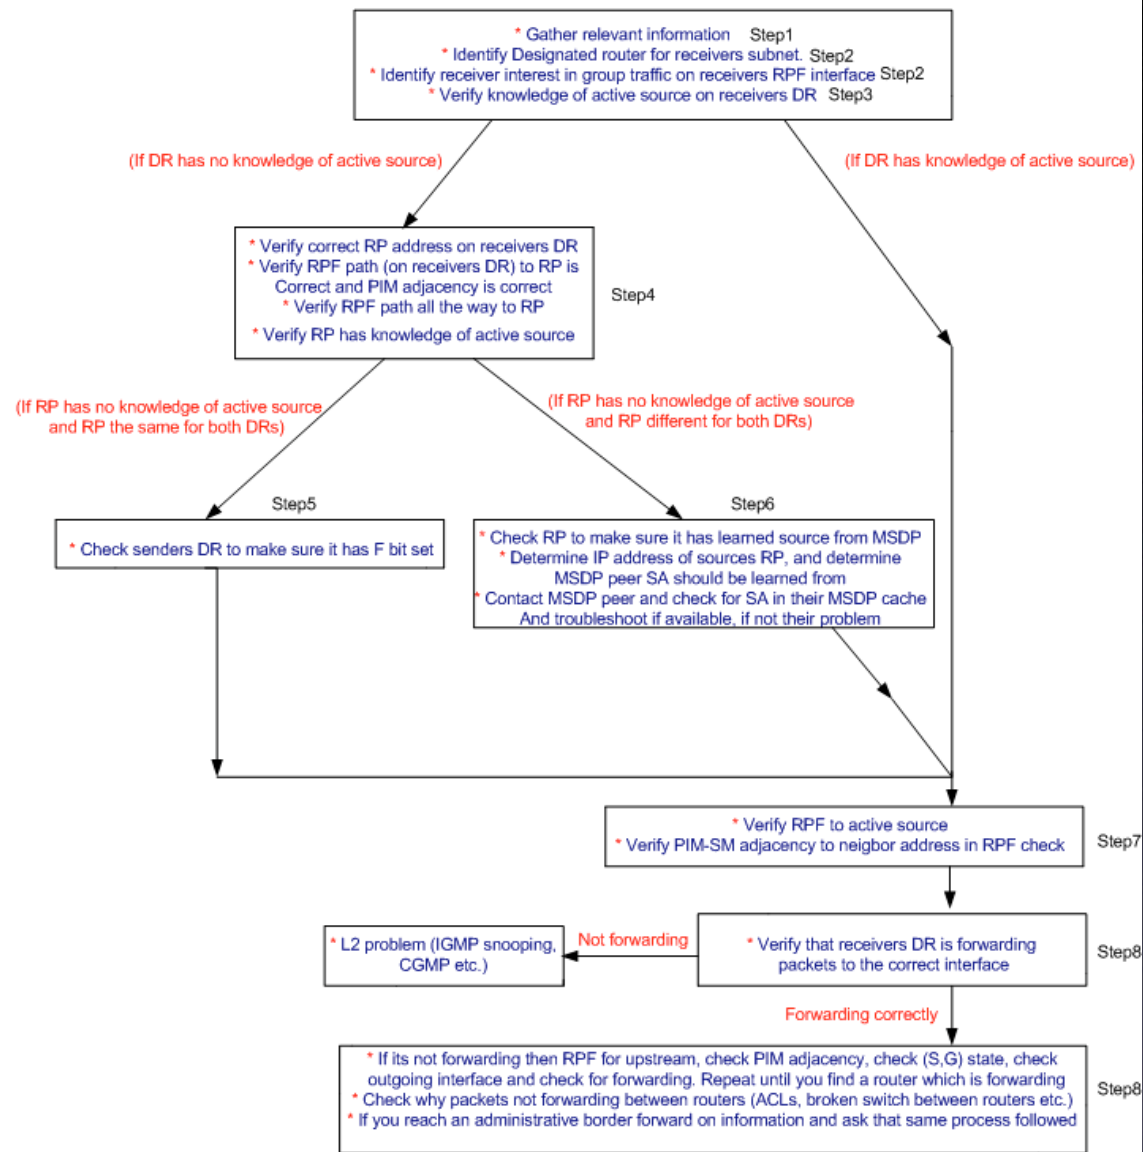
The bad ...

- ASM limitations
- Interdomain IP Multicast
- 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 on other people

Multicast troubleshooting – John Lyons (16/09/05)



Commands Required :

- | | |
|---|--|
| <ol style="list-style-type: none"> 1. Show ip route <i>receivers interface</i>
Show ip pim neighbor <i>RPF interface</i>
Show ip igmp groups <i>RPF interface</i>
Show ip mroute <i>group-address source-address</i> 2. Show ip pim rp mapping
Show ip route <i>RPF-neighbor address</i>
Show ip pim neighbor <i>RPF interface</i>
Show ip mroute <i>group-address source-address</i> | <ol style="list-style-type: none"> 3. Show ip mroute <i>group-address source-address</i> 4. Show ip mroute <i>group-address source-address</i> 5. Show ip route <i>RPF-neighbor address</i>
Show ip pim neighbor <i>RPF interface</i> 6. Show ip mroute <i>group-address source-address count</i> 7. Show ip route <i>RPF-neighbor address</i>
Show ip pim neighbor <i>RPF interface</i>
Show ip mroute <i>group-address source-address</i> |
|---|--|

The bad ...

- ASM limitations
- Inter-domain IP Multicast
- 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 access!
- Well, there are some other things you should know ...

The Ugly ...

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

The Ugly ...

- MAC address 32:1 overlapping
 - Software/hardware 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 L2 addresses
- 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
- Renumbering is never fun

The Ugly ...

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

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 software revision with 350 known 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
- Code 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
- Code 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
- Code bugs
- Vendor support
- Tunneling
- Inappropriate usage

Inappropriate usage

- Some people will try and persuade you that it's a good idea to run interdomain multicast video conferences with 20+ participants in 20+ Universities around the world
- These people obviously **do not** troubleshoot multicast problems
- A global survey of GRID users using AccessGRID showed that >20% of users constantly experienced problems with multicast
- Go figure

Conclusions

“The only thing that scares engineers more than IPv6 is multicast.”

- Multicast adoption will continue to rise
- 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 Multicast Networks
- Inter-domain IP Multicast
- Internet 2 Multicast workshops
- Litvanyi / Nickless (NANOG 27) - A methodology for troubleshooting IP Multicast
- Vendor documentation

Thank you for listening

Q & A

john@john.lyons.name