IP Multicast -
Embrace the challenge?

Paul Catchpole
Network Specialist

University of Warwick

THE UNIVERSITY OF WARWICK

Introduction

Paul Catchpole - p.catchpole@warwick.ac.uk

- What, why, who...
- Making it work
- Keeping it working
- Securing it
- What's next?

Multicast?

- What is it?
  - IPv4 Multicast
    - L3 requires mrouting proto – PIM SM/DM common
    - L2 packets include 23bit group address mapped to MAC with vendor 01:00:5e [IANA]. First octet contains multicast flag…
  - Class D addressing
    - 224/4
      - including 239/8 as site-local –
        - » But! 239.255/16 site-local to uni, rest of /8 out to JANet!
      - 224.0.1.0 – 238.255.255.255 – global
      - 232/8 – SSM reserved
      - 233/8 – Glop addressing based on AS number

Multicast?

- Why?
  - One-to-many / many-to-many
    - Saves bandwidth – check your Packeteer / Netflow / Tap for iplayer bandwidth etc! ☺
    - Network performs the replication
    - Only interested parties see the data – if you've configured it correctly…

Multicast Defined

- What for???
  - Internal uses:
    - Ghost / disk imaging
    - Internal comms / vid-conf / research
    - Digital signage / cross-campus broadcasting
  - External uses:
    - Freewire TV
    - Academic & Internet ‘streamed’ content
    - Conferencing (Access-grid et al)

Implementation thoughts!
Intradomain Routing

- Feature compatibility / operation
  - L3 support PIM-SM / PIM-SSM
    - Too few apps to go SSM-only!
  - Resilience required at RP level? MSDP?
  - IGMP (v2 and possibly v3 if SSM is reqd)
  - Mini switches / wireless
  - Client OS support
  - Addressing structure

Janet WAN Interdomain Mcast

- Address assignment
  - GLOP block based on your AS number
- Routing
  - RP – run own, or use Regional Network?
  - MSDP + MBGP
    - Change to ‘new’ BGP cli if on Cisco, you’ll need to do it anyway for IPv6
  - Check firewall PIM support
    - Eg: Cisco ASA+PIX on early 8.x code crashed on every PIM packet

Decision Time

- Implementation Decisions
  - RP choices
    - PIM BSR / (auto-rp on legacy Cisco)
    - Static assignment
    - MSDP in campus? Check underlying routing protocol support (BGP usually, and apparently newer OSPF) – Anycast RP?
  - Addressing - subnet the site-local up?
  - Check L2 switch support and miniswitches
  - Check and turn off mrouting on wlans
  - Expect software upgrades on old kit

UoW Approach - Extreme

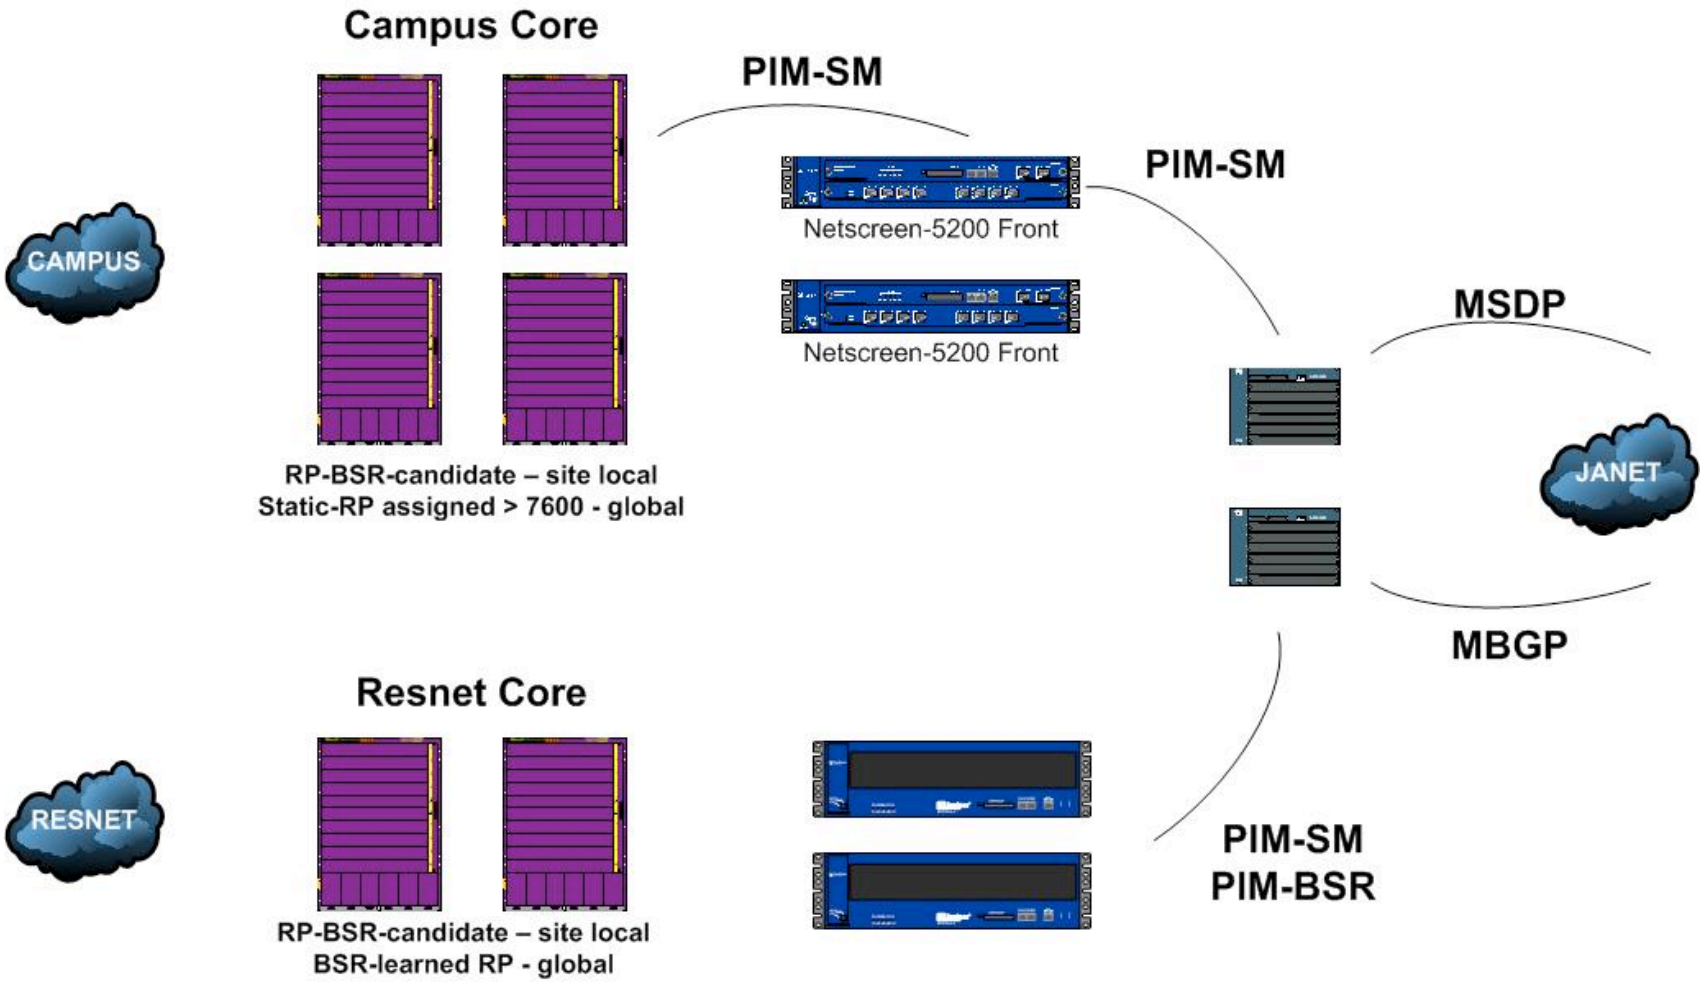

UoW Approach – Cisco NRP

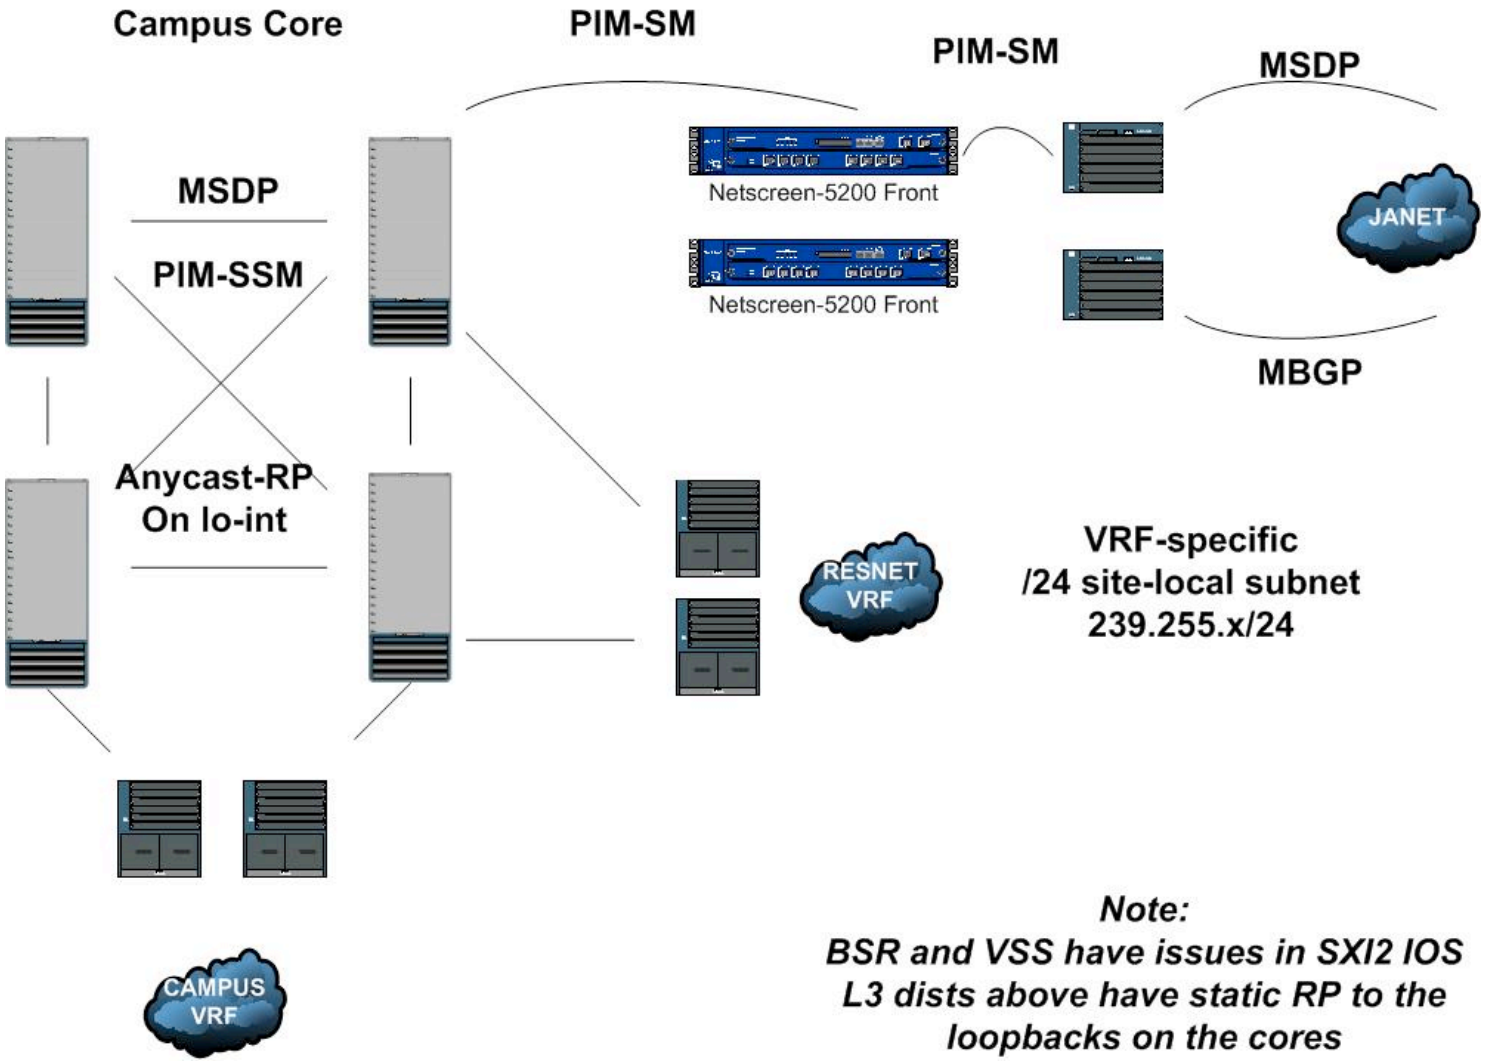

Note:
BSR and VSS have issues in SXI2 IOS
L3 dists above have static RP to the loopbacks on the cores

Common ‘Issues’

- Software support
  - First version of anything never works
  - Early Extremeware, early PIX-OS and ASA
  - IOS is exception, even early 12.x is mostly fine
  - Firewall / UTM / other-box-in-the-middle software is likely to need upgrading to recent releases
- Dumb switches
  - Lack of IGMP snooping support means flooding of packets to all ports
- Nat
  - Private-addressed receivers work
  - Private-addressed senders don’t! - typically!

More Common ‘Issues’

- Tunnelling
  - Can’t get your firewall to play ball? Use a GRE-tunnel around it, with static mroutes where needed
  - Beware though, bandwidth will be CPU-limited if box is doing it in software… [50mbit max on a 3560, probably a few hundred mbit on a VXR]
- Site local applications
  - Symantec Ghost and other similar apps often use global group addressing – be careful of rp-selection
- ‘Orrible bodges
  - Cabling round firewall or across L3 boundries with no ip addressing to get mcast working…
  - **Just say no!**

Troubleshooting

- What's seeing what where?
  - Back to basics:
    - Client on a subnet sends IGMP membership report
    - DR looks up RP for the group!
    - One DR per subnet – sends join/prune towards RP
    - Routers on path build (*,G) route
    - RP forwards sender packets 'down' the tree (usually!) until data rate exceeds configured figure – often 1 packet!
    - Downstream routers *can* join SPT from source to dest – building and showing an (S,G) entry

Troubleshooting

- Common weirdness
  - You see a (*,G) but no (S,G)
    - No packets from sender – network path issue perhaps
    - RPF check failed (won't add a folded mroute – sort of multicast split-horizon – packet duplication …)
    - No RP found or no RP reachable (unicast)
    - RP selected doesn't want to be an RP
    - ***Application TTL set to < <num L3 hops> <<< very common!***
  - Mcast flooding
    - Is IGMP Snooping turned on? Unmanaged miniswitch?

Troubleshooting

- Extreme
  - sh ipmc fdb – shows hardware m-forwarding
  - sh ipmc cache – shows mroute table, oddly…
  - sh igmp snooping <vlan> / sh igmp group <vlan>
  - sh pim rp-set – show static or learned RPs
- Cisco
  - sh ip pim rp-hash <group> - does it know the RP?
  - sh ip mroute active – what's 'live' ?
  - sh ip msdp sa-cache – is it learning group/sender/rp?
- Note: some Cisco sh-commands do not actually work in a VRF, despite allowing entry, SXI2a

Examples

- Extreme

```
  * C-SB-E-3-0A-10000:1 # sh pim rp-set
Group           Mask            C-RP            Origin    Priority
239.2.0.0       255.255.0.0     137.205.1.1     Bootstrap 0
                                137.205.1.2     Bootstrap 10

  * C-SB-E-3-0A-10000:2 # sh ipmc fdb
d_Idx   VID  Dest             Flow Source           s_Idx NxtHopIdx Ports
2ace0 (3068) 224.69.132.243   ffac 137.205.144.152  12b690 15ae0-3d1b 2:4, 8:1
3a1c2 (3068) 226.152.75.51    fff6 137.205.144.152  12b330 18f60-3f30 2:4, 8:1

  * C-SB-E-3-0A-10000:3 # sh ipmc cache
Index  Dest Group      Source              InVlan    Origin
[0000] 229.55.150.208  137.205.246.65(  ) Backb_CSB1_DUH1 PIM-SM
[0000] 224.77.129.39   137.205.1.1(R )    IN_backb PIM-SM
[0000] 224.77.129.39   137.205.195.157(  ) interlink-cisco-1 PIM-SM
[0000] 239.255.255.250 137.205.1.1(R )    IN_backb PIM-SM
```

Examples

- Cisco
  - C-SB-E-3-0A-C-7018-SW1# sh ip msdp peer vrf CAMPUS-VRF
  - MSDP peer 172.31.253.11 for VRF "CAMPUS-VRF"
  - AS 0, local address: 172.31.253.10 (loopback110)

  - C-SB-E-3-0A-C-7018-SW1# sh ip mroute vrf CAMPUS-VRF
  - (*, 239.255.255.250/32), uptime: 2w0d, pim ip
  - Incoming interface: Ethernet7/26, RPF nbr: 172.31.255.254
  - Outgoing interface list: (count: 1)
  - port-channel202.1111, uptime: 2w0d, pim

Management & Security

- Is there a problem?
  - No inherent threat from Mcast – receiver needs to ‘want’ the traffic… BUT – Consider bandwidth implications near the edge!
- Scoping
  - Filter MSDP announcements to only include valid and non-site-local groups – JANET list available
  - Multicast boundary can be used on Cisco to split organisation to allow overlapping mcast addressing (typically site-local)
- Access lists
  - Most firewalls can cope with mcast policies, access lists on Cisco and Extreme should work without problem
  - Use for usual reasons – allow/ban group traffic from entering or leaving network – filter the ‘usual’ bad application list

Management & Security

```
224.0.1.2 ! SGI- Dogfight
224.0.1.3 ! Rwhod
224.0.1.8 ! SUN NIS +
224.0.1.20 ! Any private experiment
224.0.1.22 ! SVRLOC
224.0.1.24 ! Microsoft-ds
224.0.1.25 ! nbc-pro
224.0.1.35 ! SVRLOC-DA
224.0.1.39 ! cisco-rp-announce
224.0.1.40 ! cisco-rp-discovery
224.0.1.60 ! hp-device-disc
224.0.1.76 ! IAPP – wireless base-station
224.0.2.1 ! RWHO
224.0.2.2 ! SUN RPC
224.0.2.3 ! EPSON-disc-set
224.0.23.1 ! Ricoh-device-ctrl
224.0.23.2 ! Ricoh-device-ctrl
225.1.2.3 ! Altiris
224.77.0.0/16 ! Norton Ghost
226.77.0.0/16 ! Norton Ghost
229.55.150.208 ! Norton Ghost
234.42.42.40/30 ! Imagecast disk duplication
234.142.142.142 ! Imagecast disk duplication
239.255.0.0 0.0.255.255 ! Site-local scope
```

The future...?

- Source specific multicast – SSM
  - ‘This-group-with-this-sender’ please…
  - OS Support variable
  - Backwards compatible at the router with PIM-SM
  - Requires IGMP v3 on routers and hosts
  - Uses specific group range for SSM apps

The future...

- World's quickest IPv6 mcast overview…
  - Group addresses – all within: ff00::/8
  - See RFC3306/7 for group addressing
  - Embedded-RP removes need for MSDP
  - PIM-SM works in similar way
  - IGMP function now by MLD (Listener Discovery)
  - SSM uses ff3x::/96
  - MSDP Anycast RP gone, so BSR is back, yay…

And that's it...

- Questions?
- p.catchpole@warwick.ac.uk

Well, almost…

- A question for you…
  - Cisco Cat6500 L3 SVI
    - Apply an access list
    - Which way is inbound and outbound?

- Answers on a postcard… ☺