

Multicast: Wide-Area Perspective

Jonathan J. Couzens

Principal Network Engineer

JANET Network Operations & Service Centre

jjc@nosc.ja.net

+44 (0)20-7692-1316



Agenda

- Introduction
- The Interdomain Boundary
- Addressing and Filtering

Agenda

- Introduction
- **The Interdomain Boundary**
- Addressing and Filtering

Interdomain Boundary

- What is it?
 - The network connection between two routers that are under the control of different managements
 - » connection between the SJ4 BAR and a RN's boundary router
 - » connection between a RN's boundary router and a site boundary router
 - for the purposes of illustration, we have two domains, *A* and *B*, with edge routers *A* and *B* either side of the boundary

Interdomain Boundary

- Unicast routing across the boundary
 - may use BGP4 to exchange unicast address ranges
 - router A advertises domain *A* routes to router B, and vice versa
 - may simply use static routes
 - router A has static routes for domain *B* pointing to router B, and vice versa

Interdomain Boundary

- How do we allow multicast to cross the boundary?
 - Configure each router to run PIM-SM on the interconnect
 - » allows multicast traffic to flow across the boundary
 - » but that's not enough....

Interdomain Boundary

- Sparse-mode is ‘on-request only’
 - will only deliver multicast packets if it knows that the receiver wants them
 - the receiver (strictly the last hop router) will only ask the source for traffic if it knows about it
- in a single domain, the Rendezvous-Point (RP) fulfils this function

Interdomain Boundary

- Each domain should have its own RP
 - the RP potentially controls who can source and receive what
 - it is normal to filter multicast traffic on interdomain boundaries to prevent ‘local’ groups ‘escaping’
 - » for example, multicast groups used by PC administration tools
 - if a domain uses another domain’s RP, it loses all control and is subject to the other domain’s policy

Interdomain Boundary

- Each RP only knows about sources and receivers in its own domain
 - the key is how one RP can know of another's sources
 - » if the source is known, the 'foreign' RP can join to it and initiate traffic
 - » once the receiver gets the traffic, it can join directly to the source and the RP goes out of the loop
 - the only method currently available is to use MSDP (*Multicast Source Discovery Protocol*)

Interdomain Boundary

- MSDP has had a chequered history - 1
 - It was effectively formalised by the formation of the MSDP WG of the IETF in late 1998
 - If all had gone well, an RFC detailing the specification should have been available in 1999
 - Unfortunately, it didn't
 - By the IETF meeting in July 2002, it had all but collapsed
 - 'No agenda items submitted didn't want to hold meeting'

Interdomain Boundary

- MSDP has had a chequered history - 2
 - The last act of the WG was to rationalise the existing Internet-Draft (by then, version 13)
 - The work incorporated in version 7 to 13 was effectively scrapped and version 6 was reissued (with some minor changes) as version 14
 - This has now been taken over by the MBONED WG for submission as an RFC
 - In fairness to various vendors, this is the primary reason why products have been slow to appear

Interdomain Boundary

- MSDP is all we've got
 - in due course, BGMP is projected to take over this role
 - but, until then,....

Interdomain Boundary

- MSDP is designed as an interdomain protocol
 - this means that it's really meant to operate between Autonomous Systems (ASes)
 - interAS routing uses BGP
 - hence there are dependencies on BGP in MSDP

Interdomain Boundary

- BGP can carry multiple address families
 - IPv4
 - » unicast routes
 - » multicast source routes
 - IPv6
 - VPNs

Interdomain Boundary

- BGP routes kept internally in RIBs (Routing Information Base)
 - unicast routes in the Unicast RIB (URIB)
 - multicast routes in the Multicast RIB (MRIB)
- MSDP performs its RPF checks against these RIBs, not necessarily against the router's internal forwarding table

Interdomain Boundary

- MSDP's RPF rules
 - Each Source Active (SA) announcement sent by an MSDP peer contains, as well as the Source/Group information, the RP that originated the information
 - » the RP that the source originally registered with
 - MSDP avoids loops in its topology by applying RPF checks to the originating RP address

Interdomain Boundary

- MSDP's RPF rules - 1
 - the peering is between two adjacent routers running BGP
 - MSDP searches the MRIB, then the URIB for the path to the originating RP
 - » if none, reject
 - The first AS in the path to the RP must be the same as the AS of the MSDP peer
 - » if not, reject

Interdomain Boundary

- MSDP's RPF rules - 2
 - the peering is between two non-adjacent routers running BGP (multihop eBGP)
 - MSDP searches the MRIB, then the URIB for the path to the MSDP peer address
 - » if none, reject
 - The first AS in the path to the MSDP peer must be the same as the AS of the MBGP peer
 - » if not, reject

Interdomain Boundary

- MSDP's RPF rules - 3
 - the peering is between two routers that are not running BGP
 - The normal RPF check will fail but there are ways of getting it to work:
 - » redistribute the static routes into BGP, then the RPF check will work
 - » on the downstream peer, configure the peering as a 'default peer' which will override the RPF check

Interdomain Boundary

- MSDP's RPF rules - 4
 - if there is only one MSDP peer, don't bother with the RPF check
 - » with only one peering, there can't be a loop
 - » whether BGP is running is irrelevant in this case

Agenda

- Introduction
- The Interdomain Boundary
- Addressing and Filtering

Addressing and Filtering

- Transparency
 - Certain multicast groups are used by badly thought-out applications
 - possibly appropriate for the LAN but not appropriate for the global Internet
 - » laser printers advertising themselves
 - Traffic sourced from private address ranges shouldn't be advertised

Addressing and Filtering

- Transparency
 - all multicast traffic crossing interdomain boundaries should be filtered to exclude multicast traffic belonging to certain well-known groups
 - all interdomain MSDP SA messages should be filtered for the same groups and for certain source addresses

Addressing and Filtering

- Groups that should remain local to a domain
 - 224.0.1.2 SGI-Dogfight
 - 224.0.1.3 Rwhod
 - 224.0.1.22 SVRLOC
 - 224.0.1.24 Microsoft-ds
 - 224.0.1.35 SVRLOC-DA
 - 224.0.1.60 Hp-device-discovery
 - 224.0.2.2
 - 234.42.42.42 Imagecast
 - 229.55.150.208 Norton Ghost
 - 234.142.142.142 Imagecast
 - 225.1.2.3 Altiris
 - 224.0.1.39 Auto-RP
 - 224.0.1.40 Auto-RP

Addressing and Filtering

- Administratively-scoped groups
 - 239.0.0.0 - 239.255.255.255
 - group addresses to be explicitly contained within management domains or groups of domains
 - on all interdomain boundaries, 239.255.0.0/16 should be filtered
 - JANET filters the whole of 239.0.0.0/8 on the external borders except for GEANT where we allow 239.194.0.0/16

Addressing and Filtering

- PIM-SSM groups
 - 232.0.0.0 - 232.255.255.255
 - Source-Specific Multicast
 - must be filtered out by MSDP peerings
 - RPs do not need to know about sources in this range as the receivers join to the sources directly without needing to use the RP
 - RPs must not accept registers from sources in this range

Addressing and Filtering

- GLOP groups
 - 233.0.0.0 - 233.255.255.255
 - global addresses allocated to ASes
 - does not apply to private ASes as used inside JANET
 - only 256 groups per AS
 - JANET's is 233.3.18.0/24
 - this is not currently available for general community use
 - one address in use so far
 - » 233.3.18.1 - JANET Beacon

Finally...

- Interesting reading
 - The MBONED WG of the IETF is doing a lot of work in developing standards for the operational deployment of multicast
 - draft-ietf-mboned-ipv4-mcast-bcp-00
 - draft-ietf-mboned-msdp-deploy-00
 - draft-ietf-mboned-iesg-gap-analysis-00
 - Issue 5-4 of the *Internet Protocol Journal* has a long list of references to work on future directions
 - <http://www.cisco.com/ipj>

Multicast: Wide-Area Perspective

Jonathan J. Couzens

Principal Network Engineer

JANET Network Operations & Service Centre

jjc@nosc.ja.net

+44 (0)20-7692-1316

