

CA*net 3 Multicast Services Document

Version 1.0
Nov 3, 2000

- 1: Introduction.....3**
- 2: Existence backbone configurations.....3**
 - 2.2: PIM-SM.....3
 - 2.2: MBGP4
 - 2.3: MSDP4
- 3: Multicast routing policy4**
 - 3.1: Routing Policy4
 - 3.2: Connectivity Scenarios5
 - 3.2.1: GigaPOPs running PIM-SM and supporting MBGP/MSDP.....5
 - 3.2.2: GigaPOPs running PIM-SM but not Supporting MBGP.....6
 - 3.2.3: GigaPOPs running flood and prune protocols (DVMRP/PIM-DM).....6
- 4: Multicast Peering.....7**
 - 4.1: Canadian peers.....7
 - 4.2: International Peers.....7
- 5: Administrative scope.....8**
- 6: Static Multicast Address Assignments.....8**
- Appendix A: Configuration Example for CA*net 3 Core Router.....9**

1: Introduction

Inter-domain routing, scalability, and multi-protocol negotiation pose significant challenges to ubiquitous native multicast. These problems have been most evident in our experiences with setting up, implementing, and managing multicast services on CA*net 2.

The CA*net2 backbone was designed to run PIM-SM between GigaPOPs with a single rendezvous point (RP) configured at the Toronto gateway router. DVMRP was used to exchange routing information between GigaPOPs and RANs, and PIM-DM was used to manage multicast traffic across peer links. Experience has been indicated that DVMRP and tunneling have contributed to a number of problems such as high bandwidth usage and CPU utilization.

The design goal for CA*net 3 multicast services is to take the advantage of the latest protocols for multicast routing and session propagation, and to increase the reliability of multicast services. Four key objectives were identified for CA*net 3 multicast design.

1. Eliminate all multicast tunnels on the network backbone.
2. Replace all interior routing protocols running across peer links with exterior multicast routing protocols.
3. Move away from flood and prune protocols to shared tree protocols.
4. Establish an acceptable standard for multicast peering

CA*net 3 Multicast services utilize a suite of protocols, namely PIM-SMv2 (Protocol Independent Multicast Sparse-Mode Version 2), MBGP (Multi-protocol Extension to BGP -4) and MSDP (Multicast Source Discovery Protocol), to provide reliable multicast routing.

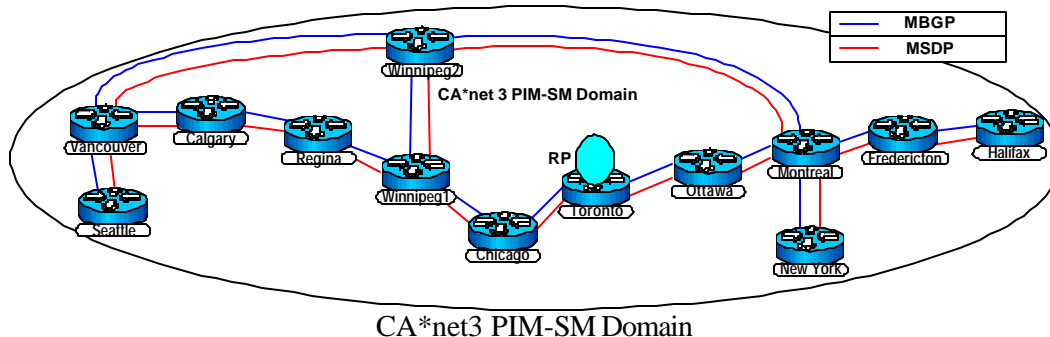
2: Existence backbone configurations

2.2: PIM-SM

The shared tree and explicit joint model of PIM-SM are well suited for a network sparsely distributed across a large area. CA*net 3 multicast routing has adopted this model. At the core, PIM-SMv2 is enabled on the backbone interfaces. A single rendezvous point (RP) is located at Toronto core at the center of the backbone. For those GigaPOPs which don't have their own PIM-SM domain, CA*net 3 domain's RP can be used as their own PIM-SM domain's RP. The command to enable this is:

```
ip pim rp-address 205.189.32.241 ( 205.189.32.241 is Toronto core router loopback address )
```

GigaPOPs should consider implementing their own PIM-SM domain, thereby allowing them to position their RP for optimal PIM-SM performance.



2.2: MBGP

MBGP (defined in RFC 2283) is an extension of the Border Gateway Protocol (BGP). It provides for the same flexibility and control as BGP does for unicast routing and peering. With MBGP, the BGP routing update contains both unicast and multicast routing information. The Multicast routing information is used to calculate the Reverse Path Forwarding (RPF) between autonomous system. Full mesh of MBGP peers is configured in CA*net 3 to provide a consistent multicast routing information base across the backbone. External MBGP is also used to peer with several international networks for exchanging multicast traffic. Multiple external MBGP peers provide redundancy in case of failure.

2.3: MSDP

MSDP is a mechanism to exchange Source-Active (SA) messages across multiple PIM-SM domains. It allows multicast source groups to be known to all rendezvous point(s) (RPs) its peer with. Each PIM-SM domain uses its own RP source information and need not depend on RPs in other domains.

In the CA*net 3 backbone, an MSDP mesh group was configured providing a consistent view of available sources within the PIM-SM domain. Any SA messages received from a peer in a mesh group are not forwarded to other peers in the same mesh group. This will reduce the SA message flooding and simplify peer-RPF flooding.

MSDP Cache Source-Active State was turned up on all backbone routers. Caching SA message reduces the join latency. Routers get source/group information from the cache instead of waiting for next SA message. By default, CISCO routers do not cache source/group pair information from SA messages.

3: Multicast routing policy

In order to maintain seamless multicast services, policies have been put in place to ensure regional traffic does not destabilize the backbone. Furthermore, reliable multicast services across CA*net 3 is dependent on the cooperation of all partners participating in multicasting. Policies for multicast connectivity can be divided into two areas: routing policy, and recommended connectivity scenarios.

3.1: Routing Policy

1. DVMRP routes from any source are disallowed in the backbone. In order to enforce this, all backbone routers have configured to drop all encountered DVMRP routes. For those RANs running DVMRP or DM protocols internally, they need to join their PIM-SM border router.
2. Tunneling is also disallowed in the backbone.
3. The minimum peering requirement with CA*net 3 core router is PIM-SM v2. Preferably GigaPOPs run their own PIM-SM domain, otherwise Toronto Core router could be used as their own Domain's RP.
4. MBGP is also preferable for inter-domain routing protocol with peers and for GigaPOPs which wish to do to do MBGP peering. The same BGP routing and tagging policies apply (Tier A and B, refer to CA*net3 routing protocol). The same route-map is applied on both unicast and multicast routing.
5. MSDP is recommended in all peering arrangements except for those GigaPOP using CA*net 3's RP as their own RP.
6. Administrative scope filter (239.0.0.0/8) is applied to all peers (see at section 5 for more details).
7. For international peers, the multicast routing and peering policies are the same as unicast. Again PIM-SM v2 is the minimum requirement and MBGP/MSDP peering is encouraged with those capable to it.

3.2: Connectivity Scenarios


Connectivity scenarios describe a number of recommended configurations for multicast participants. These scenarios describe connectivity to help GigaPOP networks receive the best possible multicast services from the backbone.

All considered there are two possible connectivity scenarios namely, GigaPOPs supporting MBGP/MSDP and those not supporting MBGP.


3.2.1: GigaPOPs running PIM-SM and supporting MBGP/MSDP

GigaPOPs supporting MBGP/MSDP can further be divided into two types::

- A. GigaPOPs running PIM-DM or DVMRP behind their PIM-SM feed to the CA*net 3 backbone will have to do a join at a DM/SM border router. This causes the border router to join all groups on the SM side so that it behaves like a DM segment. This is no different from PIM-SM/MBGP/MSDP peer.

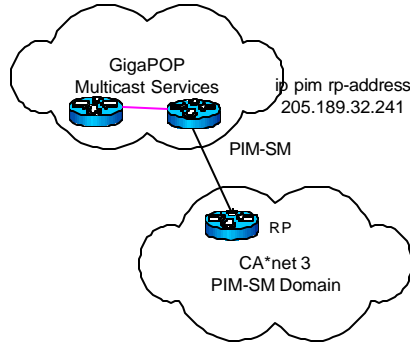


- B. PIM-SM/MBGP/MSDP connected GigaPOPs will have the same amount of control for multicast services as they do for unicast. There are no restrictions on multi-homing and no dependencies exist for inter-domain shared resources.




3.2.2: GigaPOPs running PIM-SM but not Supporting MBGP

- A. GigaPOPs without their own PIM-SM Domain, depend on the CA*net 3 backbone RP for public multicast services, should not multi-home. There are two ways to configure the RP, GigaPOP's routers either listen to RP announcement from CA*net 3's BSR or statically configure RP in the router config.



- B. GigaPOPs with their own PIM-SM Domain, but which for some reason can not do MBGP/MSDP peering, are recommended to do MSDP as the way to propagate the SA information to other PIM-SM domains, so the active multicast sources are known to other multicast networks. The interior multicast routing protocol could be DVRMP or DM protocol.



3.2.3: GigaPOPs running flood and prune protocols (DVMRP/PIM-DM)

There are no suggested methods for GigaPOPs running DVMRP/PIM-DM. The elimination of flood and prune protocols is highly recommended.

4: Multicast Peering

4.1: Canadian peers

The following table provides a current multicast view for Canadian Peers.

CA*net 3 peer	PIMv1	PIMv2	SM	MBGP	MSDP
BCnet		X	X		
NETERA		X	X	X	X
SRnet	X		X		
MRnet	X		X		
ONet	X		X		
NRC	X		X		
CRC		X	X	X	X
RISQ		X	X	X	X
UNB					
MUN	X		X		
DALnet		X	X	X	X
Wednet	X		X		
Surenet	X		X		

4.2: International Peers

PIM_SMv2 is the only allowed protocol to peer with the international networks. In order to eliminate single point of failure, CA*net 3 maintains multiple MBGP/MSDP peers with several research networks in the US and other countries. The table below shows the current international multicast peering arrangement.

International Peer	PIMv2	SM	MBGP	MSDP
VBNS	X	X	X	X
STARTAP	X	X	X	X
NORDUnet	X	X	X	X
Abilene	X	X	X	X
NREN	X	X		
CERN				
NISN				
ESnet				
ANL	X	X		
IUCC				
REUNA	X	X	X	X
DREN				
TEN-155				
PNW				

5: Administrative scope

The IANA has designated the multicast address range 239.0.0.0 to 239.255.255.255 as the administratively scoped addresses. This range of addresses can be reused within GigaPOPs' administered domains.

CA*net 3 backbone routers have set up an administratively scoped boundary on the incoming interface for multicast group addresses. A standard access list as follow is setup to deny access from either direction. The access list below also blocks the multicast group address (224.0.1.39 and 224.0.1.40) which carry auto-RP announce and discovery packets.

```
ip access-list standard multicast-boundary
deny 224.0.1.39
deny 224.0.1.40
deny 239.0.0.0 0.255.255.255
permit any
```

6: Static Multicast Address Assignments

There are occasions where require static multicast addresses are required for new tools and applications. In the past a request had to be made to IANA to receive a globally unique address assignment. Recently a new approach for allocating globally unique multicast addresses has been developed by David Meyer of Cisco Systems and Peter Lothberg of Sprint. By combining a segment of addresses from the multicast address space (233/8) with autonomous system (AS) numbers it is possible to allocate a range of globally unique multicast addresses for each AS. Static multicast allocations are discussed in detail in the RFC 2770 "GLOP addressing in 233/8".

Each CA*net 3 GigaPOP or AS administrator is responsible for a designated static allocation multicast address space to be distributed among their various subnets. The example of static multicast address space is allocated as follows:

GLOP addressing example in 233/8

GigaPOP	AS Number	Address Space
BCnet	271	233.1.15/24
NETERA	852	233.3.85/24
SRnet	803	233.3.35/24
MRnet	10965	233.42.213/24
ONet	549	233.2.37/24
CRC	818	233.3.50/24
NRC	10786	233.42.34/24
RISQ	376	233.1.120/24
UPEI	7860	233.30.180/24
DALnet	8111	233.31.175/24
MUN	10972	233.42.220/24
UNB	856	233.3.88/24
Wednet	11700	233.45.180/24
Surenet	12004	233.46.228/24

Appendix A: Configuration Example for CA*net 3 Core Router

```
!  
!Global command, enable multicast routing and set dvmrp routes limited  
!  
ip multicast-routing distributed  
ip sdr cache-timeout 20  
ip dvmrp route-limit 20000  
!  
! Core Routers interface configure, enable PIM-SM v2 on intre-connected interfaces  
!  
interface POS0/0  
description OC48 - Tor to Chi link, NMC facility ID:  
ip address 205.189.32.146 255.255.255.252  
no ip directed-broadcast  
ip pim sparse-mode  
ip router isis  
ip mroute-cache distributed  
ip sdr listen  
!  
! Define RP location, and default mroute to Toronto core  
!  
ip pim rp-address 205.189.32.241  
ip pim bsr-candidate Loopback0 24 255  
ip mroute 0.0.0.0 0.0.0.0 205.189.33.166 250  
  
!  
! BGP/MBGP config, meshed of BGP and MBGP at core routers  
!  
neighbor 205.189.32.242 remote-as 6509 nlri unicast multicast  
neighbor 205.189.32.242 update-source Loopback0  
neighbor 205.189.32.242 send-community  
neighbor 205.189.32.242 version 4  
neighbor 205.189.32.242 soft-reconfiguration inbound  
neighbor 205.189.32.243 remote-as 6509 nlri unicast multicast  
neighbor 205.189.32.243 update-source Loopback0  
neighbor 205.189.32.243 send-community  
neighbor 205.189.32.243 version 4  
neighbor 205.189.32.243 soft-reconfiguration inbound  
.  
.  
!  
! MSDP config, meshed of MSDP at the core router and enable sa-cache(disable by default).  
!  
ip msdp peer 205.189.32.254 connect-source Loopback0  
ip msdp peer 205.189.32.253 connect-source Loopback0  
.  
.  
ip msdp mesh-group c3 205.189.32.254  
ip msdp mesh-group c3 205.189.32.253  
.  
.  
ip msdp cache-sa-state  
!  
! PIM-SM v2 peer, enable PIM-SM, setup administrative scope and prevent BSR message propgation.
```

```
!  
interface GigabitEthernet2/0.1  
description link to NETERA Calgary  
encapsulation dot1Q 2  
ip address 205.189.32.198 255.255.255.252  
no ip directed-broadcast  
ip pim bsr-border  
ip pim sparse-mode  
ip multicast boundary multicast-boundary  
no cdp enable  
!  
! BGP/MBGP peering config, same route-map apply on both unicast and multicast routes  
!  
neighbor 205.189.32.57 remote-as 852 nlri unicast multicast  
neighbor 205.189.32.57 next-hop-self  
neighbor 205.189.32.57 send-community  
neighbor 205.189.32.57 remove-private-AS  
neighbor 205.189.32.57 version 4  
neighbor 205.189.32.57 soft-reconfiguration inbound  
neighbor 205.189.32.57 route-map netera_in in  
neighbor 205.189.32.57 route-map netera_out out  
!  
! MSDP external peer  
!  
ip msdp peer 205.189.32.57  
ip msdp description 205.189.32.57 Netera MSDP Peer  
ip msdp sa-filter in 205.189.32.57 list 105  
ip msdp sa-filter out 205.189.32.57 list 105  
!  
! SA-filter, deny private IP, administrative scope address and vary protocol update.  
!  
access-list 105 deny ip any host 224.0.1.2  
access-list 105 deny ip any host 224.0.1.3  
access-list 105 deny ip any host 224.0.1.22  
access-list 105 deny ip any host 224.0.1.24  
access-list 105 deny ip any host 224.0.1.35  
access-list 105 deny ip any host 224.0.1.39  
access-list 105 deny ip any host 224.0.1.40  
access-list 105 deny ip any host 224.0.1.60  
access-list 105 deny ip any host 224.0.2.2  
access-list 105 deny ip any 239.0.0.0 0.255.255.255  
access-list 105 deny ip 10.0.0.0 0.255.255.255 any  
access-list 105 deny ip 127.0.0.0 0.255.255.255 any  
access-list 105 deny ip 172.16.0.0 0.15.255.255 any  
access-list 105 deny ip 192.168.0.0 0.0.255.255 any  
access-list 105 permit ip any any
```