


# Multicast Addresses

Piyush Harsh

CONS Weekly Presentation  
April 11 2007




# How it all started?



- Steve Deering worked on a project on Distributed OS called “Vsystem”
- Computers in “Vsystem” could send messages to group of different computers using Ethernet multicasting
- As project progressed, bunch of computers were added that were on other side of the campus connected via a production router
- Task of extending the MAC layer multicasting over to layer 3 fell on Steve Deering
- He developed IGMP as part of his PhD thesis and some years later developed DVMRP.

# \$16,000 caused host inefficiency



# Multicast IP Address to Multicast Ethernet Address

- For a Multicast IP packet to reach the group member on LAN, the IP packet must be encapsulated in layer 2 Ethernet frames
- Ethernet Addresses are 48 bits of which 24 bits OUI is distributed by IEEE, rest 24 bits can be varied.
- IP Multicast Addresses have first 4 bits fixed as 1110
- IP Multicast Addresses:  
1110xxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx
- Hence 28 bits in IP address can be varied

# contd...

- MAC Address assigned for IP Multicast is only 23 bits long!  
How come?
- Deering's advisor could apply for only 1 OUI from IEEE for \$1000 instead of 16 OUI Deering desired.
- Further his advisor was kind enough to give half of the addresses to play with (only 23 bits)
- The OUI for IP Multicast Address is: 01-00-5E (hex)
- The remaining 24 bits can vary from 00-00-00 to 7F-FF-FF (first bit of the 24 variable bits is 0)
- Hence 28 IP bits have to map onto 23 MAC bits

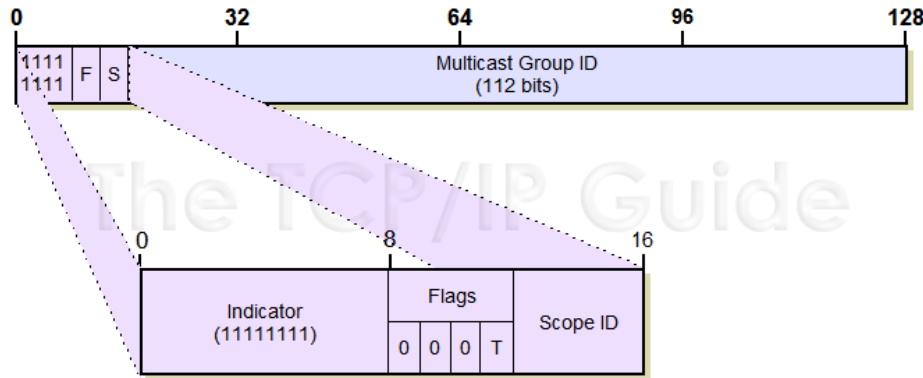


# Contd ...

- Lower 23 bits of IP Multicast group address are copied as is into the 23 bits of MAC Multicast address (higher 5 bits are ignored)
- Thus we have 32:1 ambiguous mapping which leads to over subscription problem at hosts.
- 32 different IP Multicast Addresses maps to same MAC Address, IP Stack in end hosts must take care of this over subscription. This can affect host's available CPU power if “spurious” group traffic is high enough!
- If only they had \$16,000 ...



# M-Cast Address: Scarce Commodity?



- ▶ Not if IPv6 is used.
- ▶ Problem arises in IPv4
- ▶ CIDR Notation: 224.0.0.0/4
- ▶ Seems plenty of addresses but not if every Multicast source starts using a unique IP address
- ▶ Allocation of Multicast addresses strictly controlled by IANA (Internet Assigned Numbers Authority)

## Multicast Addresses:



Class D IP Addresses are used for multicast.

# IANA Multicast Address Assignments

RANGE	Mask	Description
224.0.0.0 – 224.0.0.255	224.0.0/24	Local Network Control Block
224.0.1.0 – 224.0.1.255	224.0.1/24	Internetwork Control Block
224.0.2.0 – 224.0.255.255	-	Ad Hoc Block
224.1.0.0 – 224.1.255.255	-	Unassigned
224.2.0.0 – 224.2.255.255	224.2/16	SDP/SAP Block
224.3.0.0 – 231.255.255.255	-	Unassigned
232.0.0.0 – 232.255.255.255	232/8	Source Specific Multicast Block (SSM)
233.0.0.0 – 233.255.255.255	233/8	GLOP Block
234.0.0.0 – 238.255.255.255	-	Unassigned
239.0.0.0 – 239.255.255.255	239/8	Administratively Scoped Block

- ☆ Of these address blocks, IANA assigns addresses in the Local Network Control, Internetwork Control and Ad Hoc block based on the guidelines specified by IETF
- ☆ Addresses assigned by IANA and in the GLOP Block is static in nature
- ☆ Addresses in all other block can be allocated dynamically except for scope relative addresses whose purpose and assignments are again controlled by IANA.



# Local Network Control Block

- Range 224.0.0.0 – 224.0.0.255 is considered Link-Local Block
- Packets with address in this range are local in range
- TTL should always be 1 so that they don't go beyond local subnet

Address	Usage
224.0.0.1	All Hosts
224.0.0.2	All Multicast Routers
224.0.0.3	Unassigned
224.0.0.4	DVMRP Routers
224.0.0.22	IGMP


# Important...

- CISCO Proprietary protocol CGMP (CISCO Group Management Protocol) along with IGMP Snooping by smart CISCO Ethernet switches reduces unnecessary port flooding on LAN ports
- But Multicast addresses in Link-Local range are flooded on all LAN ports anyway because otherwise some protocols like OSPF would break
- 224.0.0.0/24 maps to 0x0100.5E00.00XX MAC Range
- There exists a 32:1 address mapping ambiguity as explained earlier
- It is therefore recommended that Multicast addresses in other class that maps to 0x0100.5E00.00XX MAC Address should be avoided

# Inter-Network Control Block

- 224.0.1.0 – 224.0.1.255 is Inter-Network Control Block
- Similar to Link-Local Block except these are used by network protocols
- Used when control messages need to be Multicast beyond the local network segment

Address	Use
224.0.1.1	Network Time Protocol
224.0.1.21	DVMRP on MOSPF
224.0.1.32	Mtrace
224.0.1.39	Cisco-rp-announce
224.0.1.40	Cisco-rp-discovery



# Ad Hoc Multicast Block

- 224.0.2.0 – 224.0.255.255 is Ad-Hoc block
- Applications which do not fall clearly into either Local network Control and Inter-Network Control categories have been assigned addresses from this block by IANA
- Most of these assignments were made by IANA prior to receiving clear guidelines from IETF, some examples below ...

Address Range	Usage
224.0.15.0 – 224.0.15.255	Hewlett Packard
224.0.12.0 – 224.0.12.63	Microsoft and MSNBC
224.0.19.0 – 224.0.19.63	Walt Disney Company

# GLOP Multicast Block

- Statically assigned experimental address range by IANA
- Addresses in this block are set by content providers and ISPs who already have a set Autonomous System Number and not by IANA
- GLOP provides for 255 Multicast addresses per ASN
- Format for addressing: 233.A.B.0 - 233.A.B.255 where A.B is the dotted decimal notation of any ASN (Autonomous System Number)
- These addresses can be leased from ISPs without intervention by IANA

```
+0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+++++
|-----233-----|-----16 bits AS-----|--local bits---|
+++++
```

# Multicast Scoping?

- Scoping is the technique which determines how far (router hops) a Multicast packet is allowed to travel before it is dropped
- Two Scoping Technique in use:
  - TTL Scoping
  - Administrative Scoping

# TTL Scoping

- The TTL value of the Multicast packet is used to determine its scope, border routers have TTL Threshold and if a arriving Multicast traffic TTL value is below the threshold, it is dropped
- Inefficient because different routers have different hop distance from a source and hence TTL scoping of Multicast packets may lead to traffic leak beyond the intended scope if configured improperly
- Traditional TTL values and their scopes
  - 0 : node local
  - 1 : Link Local
  - Up to 32 : Site Local
  - Up to 64 : Region Local
  - Up to 128 : Continent Local
  - Up to 255 : Global

# Administrative Scoping

- 239.0.0.0 – 239.255.255.255 address block reserved as Administratively Scoped addresses
- Primarily for use in private Multicast domain
- Similar in nature to reserved IP Unicast range 10.0.0.0/8
- Boundary routers must be configured to drop any Multicast packets belonging in this address range from entering/leaving the domain
- Helps in conserving limited Multicast address space
- These addresses can be reused across different domains



# Important ...

- Keep in mind that IANA reserves and defines usage for highest 255 addresses in any Administrative Scoped address range
- Used by protocols within a scope in order to function properly
- Scope Relative addresses use offset address assignments

Relative	Description
0	SAP (Session Announcement Protocol)
1	MADCAP Protocol (DHCP for Multicast)
2	SLPv2 (Server Location Protocol) Discovery
3	MZAP
4	Multicast Discovery of DNS Services
5	SSDP
6	DHCP v 4
7	AAP (Authoritative Access Point)
8	MBUS (Maintenance bus)
9-252	Reserved – To be assigned by IANA
253	Reserved
254-255	Reserved – To be assigned by IANA

# Important ... Contd .

Example:

- Consider Administratively Scoped Address block 239.146.128.0/17 for your network!
- 11101111.10010010.1xxxxxxx.xxxxxxxx
- How many addresses do you believe are available to you for use?
- 128 x 256 ? NO!
- 239.146.255.0 – 239.146.255.255 are reserved by IANA
- Available Addresses for use:  $(128 - 1) \times 256$

# Scope Relative Address Contd...

Question?

So what will be the Multicast address for the MADCAP server in the example domain above??

– 239.146.255.254

Question?

What service should be running at this address??

239.146.255.249

– DHCP v4 Server

Relative	Description
0	SAP (Session Announcement Protocol)
1	MADCAP Protocol (DHCP for Multicast)
2	SLPv2 (Server Location Protocol) Discovery
3	MZAP
4	Multicast Discovery of DNS Services
5	SSDP
6	DHCP v 4
7	AAP (Authoritative Access Point)
8	MBUS (Maintenance bus)
9-252	Reserved – To be assigned by IANA
253	Reserved
254-255	Reserved – To be assigned by IANA

# Address Allocation Techniques

- Static Address Allocation
  - Global: IANA Assignments, GLOP
  - Enterprise: Internal Assignments
- Scope Relative Address Allocation
  - Global: IANA Assignments
  - Enterprise: IANA Assignments
- Dynamic Address Allocation
  - Global: SDR, Multicast Address Set Claim (MASC), SSM (IGMP v3)
  - Enterprise: MADCAP (Similar to DHCP), SSM

# References

- ★ Guidelines for Enterprise IP Multicast Addressing, CISCO White Paper
- ★ java.net Class Inet4Address Documentation: [java.sun.com](http://java.sun.com)
- ★ Wikipedia: [http://en.wikipedia.org/wiki/Multicast\\_address](http://en.wikipedia.org/wiki/Multicast_address)
- ★ Interdomain Multicast Routing, ISBN 0-201-74612-3
- ★ Developing IP Multicast Networks, Volume 1, ISBN 1-57870-077-9

