IPv6 Multicast

St. Louis Unix Users Group
Presented March 8, 2017
By David Forrest

David Forrest graduated from Oregon State University in Business Administration, Finance Emphasis, Physics & Mathematics. Lifelong hobbyist in IT from IBM 1401, Model 20, Model 30, Model 85, Sigma 7, XDS 7, SWTP 6800, M6809, 8080, 80286, 80386, OS2, and on to currently running XP, CentOS6/7, Raspbian, Mint, and Chrome on various local and cloud machines.

Multicasting

The transmission of a packet to multiple destinations in a single send operation

IPv6 does not implement traditional IP broadcast. In IPv6, the same result can be achieved by sending a packet to the link-local all nodes multicast group at address ff02::1,

IPv6 multicast addressing shares common features and protocols with IPv4 multicast. And has new multicast implementations, including embedding rendezvous point addresses in an IPv6 multicast group address.

Routable!

In IPv4 it is very difficult for an organization to get even one globally routable multicast group assignment, and the implementation of inter-domain solutions is arcane. Unicast address assignments by a local Internet registry for IPv6 have at least a 64-bit routing prefix, yielding the smallest subnet size available in IPv6 (also 64 bits). With such an assignment it is possible to embed the unicast address prefix into the IPv6 multicast address format, while still providing a 32-bit block, the least significant bits of the address, or approximately 4.2 billion multicast group identifiers. Thus each user of an IPv6 subnet automatically has available a set of globally routable source-specific multicast groups for multicast applications

Trees!

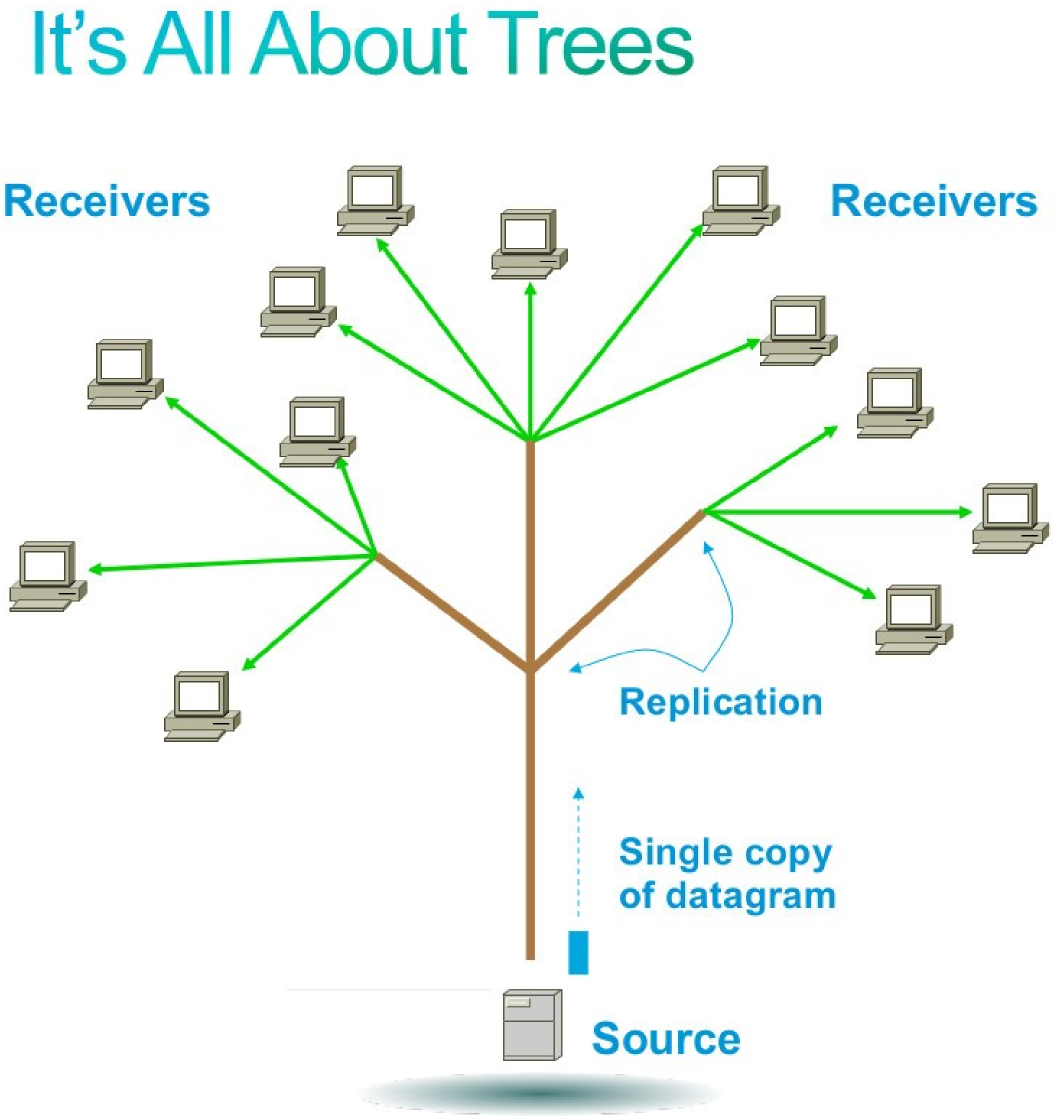
It’s All About Trees

Protocol Independent Multicast (PIM)

Link Operations - Routing Protocol - Distance Learning Surveillance – Metering - Broadcast Video Service - all with Efficient Delivery!

Including Applications We Haven't Even Built Yet

Large Privately Owned Multicast Address Space

Built-in Scoping

No NAT required, Embedded RP, Anycast, Etc..

Multicast is Foundational in IPv6

IPv6 – The Future Is NOW!

IPv6

In IPv6, there’s no longer any broadcast – sending one packet to a large number of IPv6 unspecified hosts.

There’s only multicast, unicast, and anycast. In IPv6 all nodes are required to support multicast. Without multicast, many services that you need will simply not work.

There’s a constant hum from the network on the multicast channels.

IPv6 Multicast

Multicast: An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to a multicast address is delivered to all interfaces identified by that address, either by the protocol definition or by the interface requests being joined.

There are no broadcast addresses in IPv6 as their function is superseded by multicast addresses.

Ff00::/8

1111 1111 – the code for Multicast Datagrams. As such are NOT routed on the public internet, the remaining 120 bits do not contain the normal prefix or addressing schema but do have scoping and other data necessary to provide activities per the data.

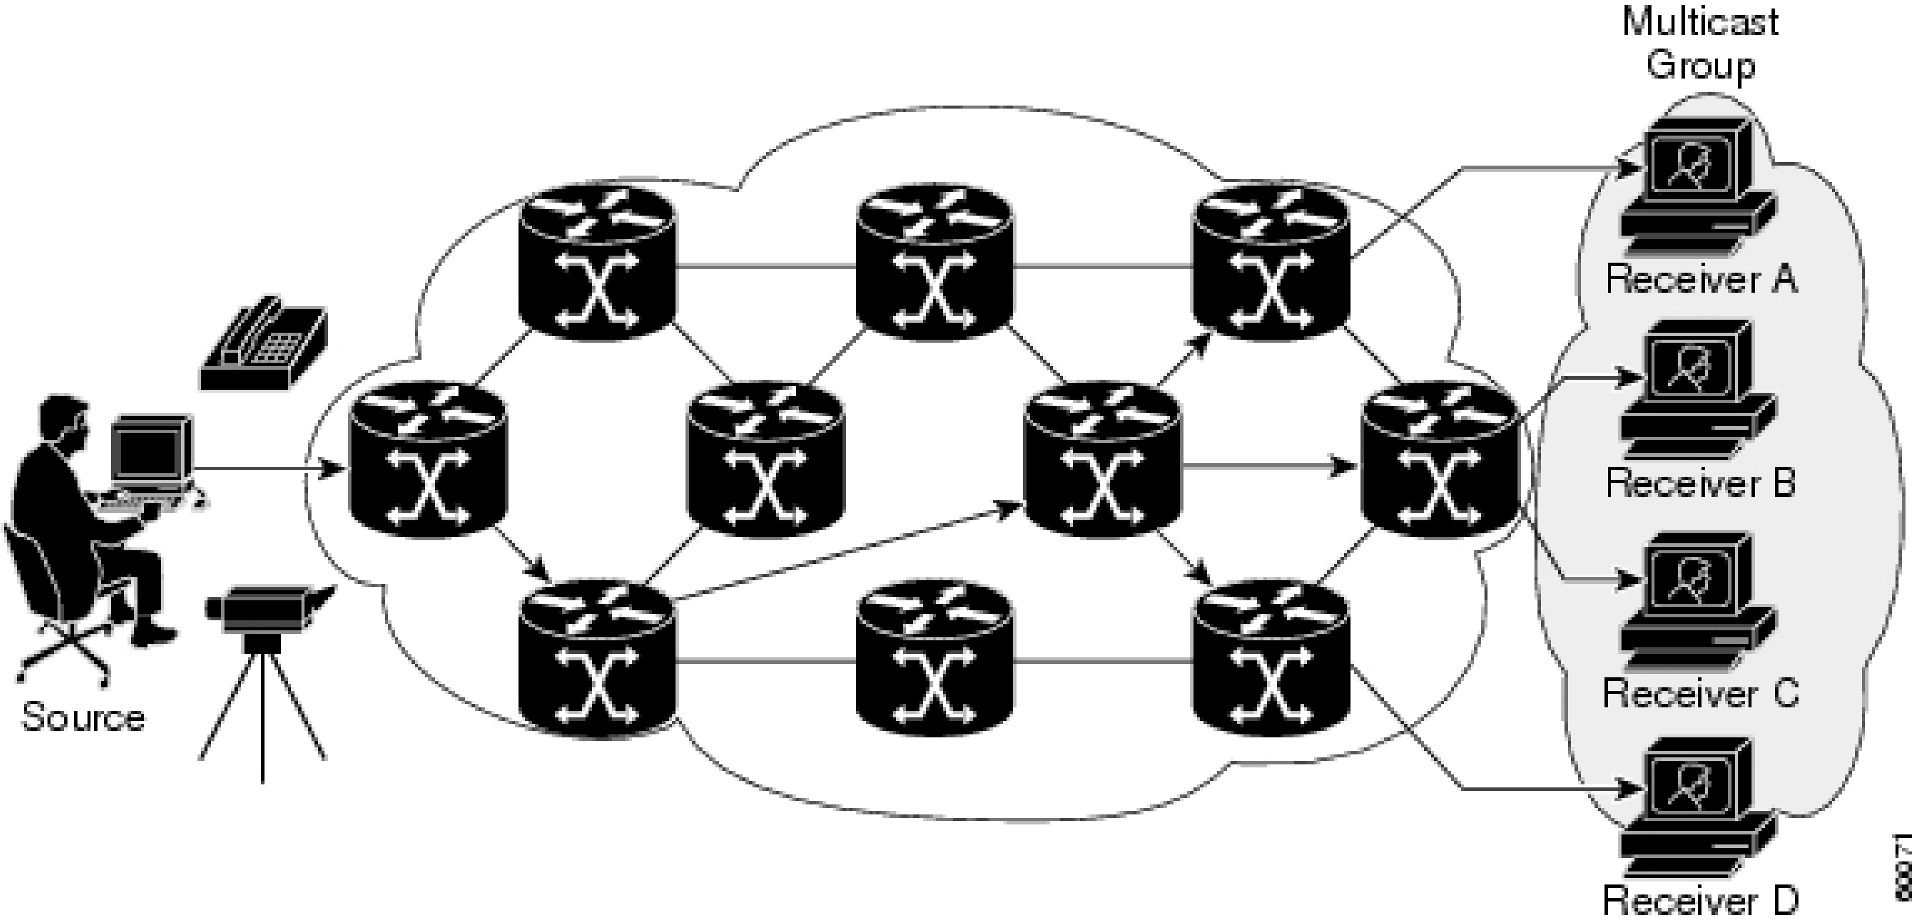
Multicast Group
Source
Receiver A
Receiver B
Receiver C
Receiver D
86071

One to many (S,G)

- The simplest form of a multicast distribution tree is a source tree with its root at the source and branches forming a spanning tree through the network to the receivers.
- Because this tree uses the shortest path through the network, it is also referred to as a shortest path tree (SPT).

SPT (S,G)

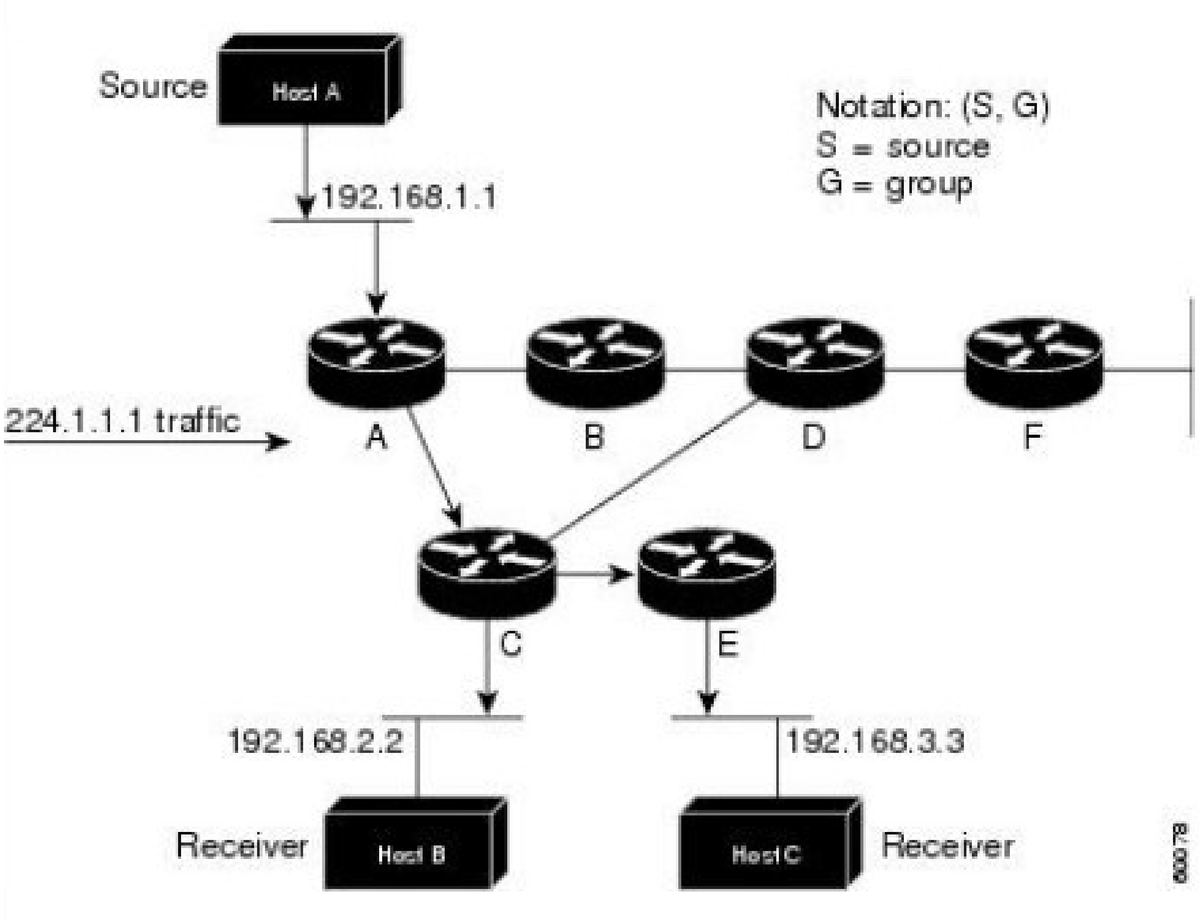
Source
Host A
192.168.1.1
Notation: (S, G)
S = source
G = group
224.1.1.1 traffic
A
B
D
F
C
E
192.168.2.2
192.168.3.3
Receiver
Host B
Host C
Receiver
60078

Many to Many (*, G)

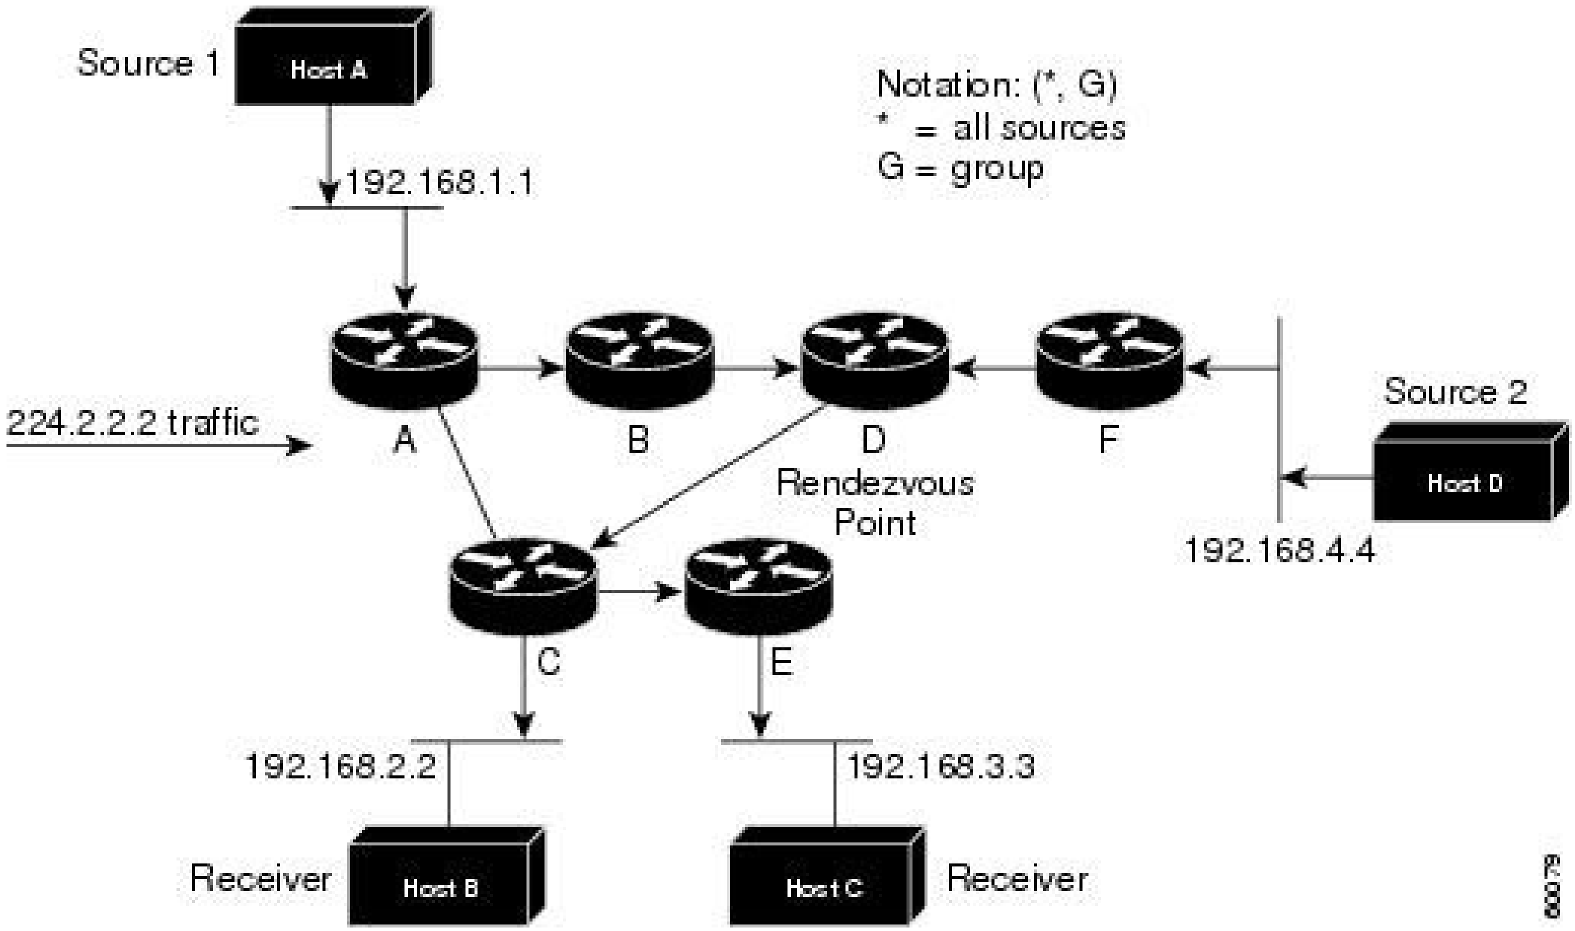

1111 1111 0RPT Scope ::/16

R = 0 No embedded rendezvous point, 1= Embedded RP

P = 0 Not unicast based, 1= Based on unicast

T = 0 Permanent address (IANA assigned)

T = 1 Temporary address (local assigned)

Scope is four bits leaving 112 bits for group ID

1=Node, 2=Link, 3=Subnet, 4=Admin, 5=Site,
8=Organization, and E=Global

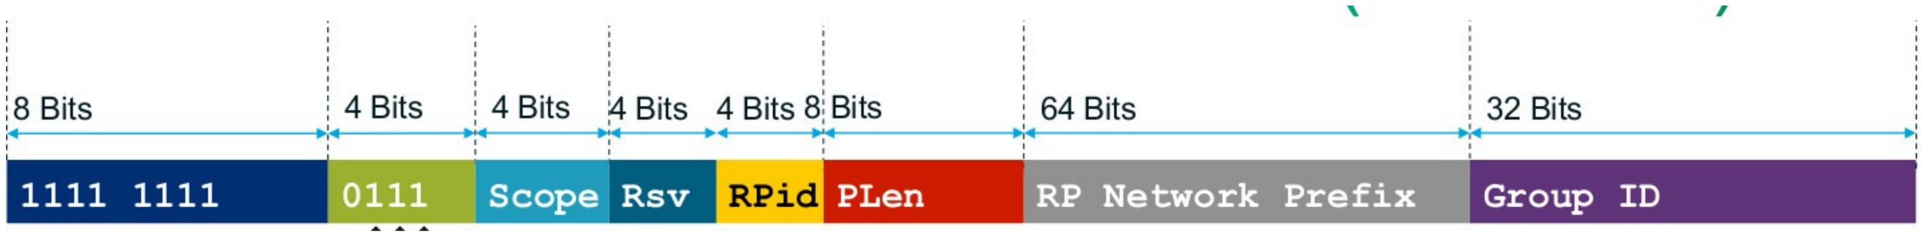

I'm going to skip describing all the fields here

Has an embedded rendezvous point;
Based on unicast;
Temporary;
The scope, rendezvous point prefix, and group ID.

Embedded RP Example

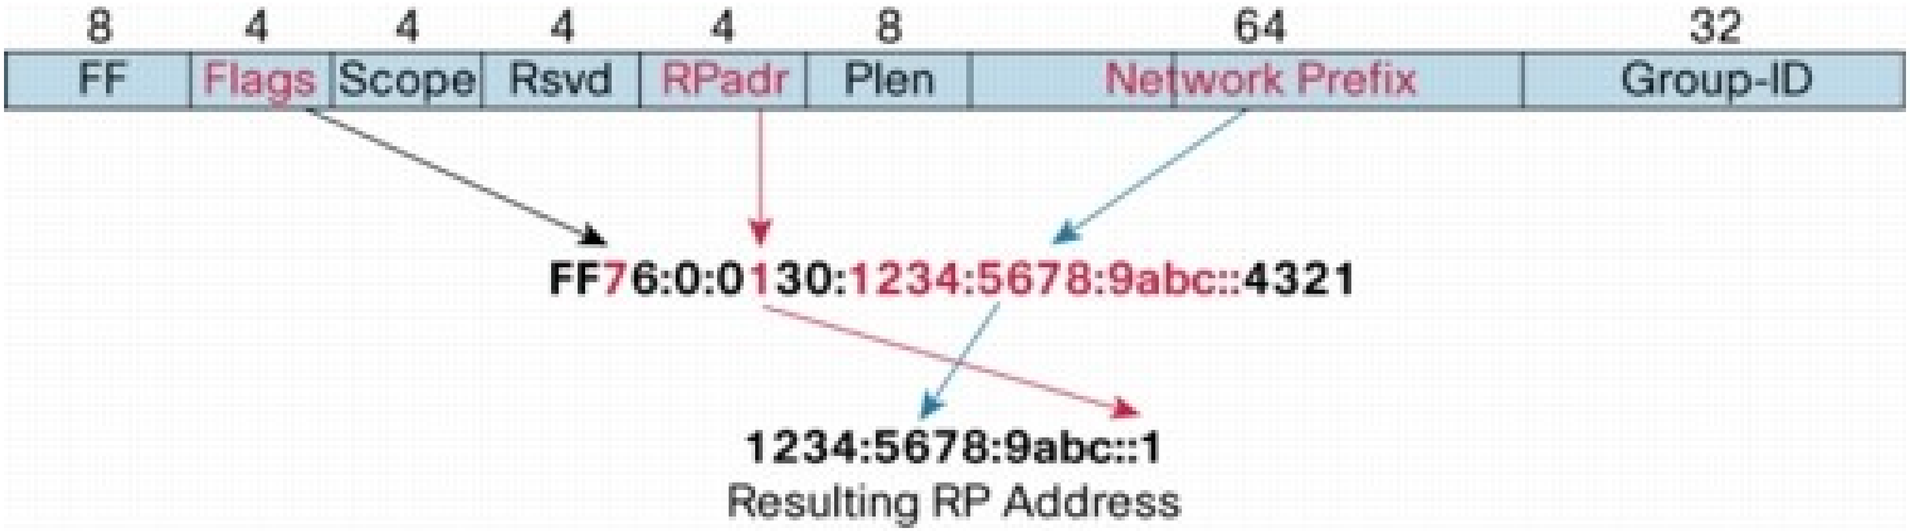

Well Known Multicast Addresses

Address	Scope	Meaning
FF01::1	Node-Local	All Nodes
FF05::2	Site-Local	All Routers
FF02::1	**Link-Local**	**All Nodes**
FF02::2	**Link-Local**	**All Routers**
FF02::A	Link-Local	EIGRP
FF02::C	Link-Local	SSDP – MSFT
FF02::FB	Link-Local	MDNS - Apple

- FF02, is a permanent address and has link scope
- Rather “Chatty” and running in your network now!

IPv6 Multicast Address	Type	RFC or Other Reference	Notes
ff02::1	ICMPv6	4291	IANA Link-Local All Nodes Address Router Advertisements are sent to this address, e.g. ICMPv6 Option (Prefix information), Sample 64-bit prefix: 2001:db8:2::
ff02::1:2	DHCPv6	3315	All DHCP Agents (servers or relays), Link-Local Scope On Windows 7 LANs, the Option Request in packets sent to ff02::1:2 is always for one of: 1. Domain Search List (24) 2. DNS recursive name server (23) 3. Vendor-specific Information (17) 4. FQDN (39)
ff02::1:3	LLMNR	4795	Microsoft’s counterpart to Apple’s mDNS Used for link-local name resolution where a central DNS server either isn’t available or appropriate.
ff02::2	ICMPv6 Router Solicitation	4291	ICMPv6 type 133 (Router Solicitation) PCs looking for an IPv6 router send packets to this address

ff02::16	ICMPv6 MLDv2	3810	ff02::16 is the Link-Local scope all-MLDv2 routers address. A packet sent to this address reaches all MLDv2 routers on a sub-network. Packets are always of Type = 143 (Multicast Listener Report Message v2) of the form Mode = Include or Exclude
ff02::c	WS-Discovery, SSDP	SSDP v1.03 WS-Discovery Specification	Universal Plug and Play (UPnP) and other Windows Rally technologies.
ff02::fb	mDNS	DNS-based Service Discovery	Multicast DNS (mDNS), Apple’s counterpart to Microsoft’s LLMNR
Solicited Node, source address is ::	Neighbor Solitication	2461	Solicited Node addresses always begin with ff02::1:ff and end with final 6 octets of proposed source address, e.g. ff02::1::ff45:3042, ff02::1::ff99:c1cc, etc.

My system

```
[drf@dave ~]$ ip maddr show dev br0
3:br0
link  33:33:00:00:00:01
link  01:00:5e:00:00:01
link  33:33:ff:7f:ad:78
link  01:00:5e:00:00:fb
link  33:33:00:00:00:fb
inet  224.0.0.251 users 2   (virbr0 & docker0)
inet  224.0.0.1
inet6 ff02::fb
inet6 ff02::1:ff7f:ad78 users 3  (enp3s0, br0, ?docker0 Via EB? )
inet6 ff02::1
inet6 ff01::1
[drf@dave ~]$
```

/sbin/ebtables -Ln

- Bridge chain: INPUT, entries: 1, policy: ACCEPT -j INPUT_direct
- Bridge chain: FORWARD, entries: 1, policy: ACCEPT -j FORWARD_direct
- Bridge chain: OUTPUT, entries: 1, policy: ACCEPT -j OUTPUT_direct
- Bridge chain: INPUT_direct, entries: 0, policy: RETURN
- Bridge chain: OUTPUT_direct, entries: 0, policy: RETURN
- Bridge chain: FORWARD_direct, entries: 0, policy: RETURN

My loopback interface

1: lo

inet 224.0.0.1

inet6 ff02::1

inet6 ff01::1

My wired interface

```
2: enp3s0
link  33:33:00:00:00:01
link  01:00:5e:00:00:01
link  33:33:ff:7f:ad:78
inet  224.0.0.1
inet6 ff02::1:ff7f:ad78
inet6 ff02::1
inet6 ff01::1
```

A collateral virtual

```
4: ip_vti0
inet6 ff02::1
inet6 ff01::1
```

My virbr0

```
5: virbr0
link  01:00:5e:00:00:01
link  01:00:5e:00:00:fb
inet  224.0.0.251
inet  224.0.0.1
inet6 ff02::1
inet6 ff01::1
```

A virtual bridge nic

6: virbr0-nic

inet6 ff02::1

inet6 ff01::1

docker0

```
7: docker0
link  33:33:00:00:00:01
link  01:00:5e:00:00:01
link  01:00:5e:00:00:fb
inet  224.0.0.251
inet  224.0.0.1
inet6 ff02::1
inet6 ff01::1
```

Kernel /proc listing

```
[root@dave ~]6 cat /proc/net/if_inet6
00000000000000000000000000000001 01 80 10 80       lo
fe80000000000000e23f49fffe7fad78 03 40 20 80      br0
fe80000000000000e23f49fffe7fad78 02 40 20 80  enp3s0
20014978000f8640e23f49fffe7fad78 03 80 00 80      br0
fe800000000000000042f3fffe0a2b3c 07 40 20 80  docker0
2602030631a84e40e23f49fffe7fad78 03 40 00 00      br0
[root@dave ~]7
```

Multicast Source Discovery Protocol (MSDP)

- MSDP was developed for peering between Internet service providers (ISPs). ISPs did not want to rely on an RP maintained by a competing ISP to provide service to their customers. MSDP allows each ISP to have its own local RP and still forward and receive multicast traffic to the Internet.

  Explaining this protocol is beyond the scope of this hobbyist! Give thanks that others can.

Source Specific Multicast (SSM)

SSM is an extension of the PIM protocol that allows for an efficient data delivery mechanism in one-to-many communications. SSM enables a receiving client, once it has learned about a particular multicast source through a directory service, to then receive content directly from the source, rather than receiving it using a shared RP.

SSM removes the requirement of MSDP to discover the active sources in other PIM domains. An out-of-band service at the application level, such as a web server, can perform source discovery. It also removes the requirement for an RP.

Bidirectional PIM

- Deploying Bidirectional PIM for Many-to-Many Applications
- This document attempts to provide self-standing guidelines on the deployment of bidirectional Protocol Independent Multicast (PIM), and includes an introduction to the protocol, design and configuration guidelines for a successful deployment, and some case studies of implementations.
- Bidirectional PIM. Bidirectional PIM is a member of the suite of multicast routing protocols supported in Cisco IOS® Software. The family of PIM protocols includes dense-mode, sparse-mode, source specific.multicast (SSM), and bidirectional (Bidir) PIM. The initial set of protocols only included dense-mode and sparse-mode, but after a few years of deployment experience, the protocols have evolved and been optimized to better support the emerging multicast applications. Mainly, SSM was developed to easily support one-to-many applications by greatly simplifying the protocol mechanics for deployment ease. On the other hand, Bidir PIM was developed to help deploy emerging communication and financial applications that rely on a many-to-many applications model. Bidir PIM enables these applications by allowing them to easily scale to a very large number of groups and sources by eliminating the maintenance of source state.
- As you can see, these two protocols, SSM and Bidir, serve both ends of the spectrum of multicast applications: one-to-many and manyto-many. For this reason, neither one is a replacement for sparse-mode PIM, which is able to support the full spectrum of multicast applications but with a bit more complexity and overhead. Given the characteristics of the different variations of the protocols, they are meant to work side by side in the network. The network operator is now able to simultaneously deploy SSM for corporate communication applications, Bidir for “hoot-n-holler” and financial applications, and sparse-mode PIM for general IP Multicast connectivity. Later in this document, we will discuss how the IP Multicast address range is utilized to support concurrent deployment of the various PIM modes.

Take aways

- Some of the following comments are opinion.
- Please accept my appreciation for your attention to the possibilities IPv6 has in the developing environment. I was one of the early Hurricane Electric certified “Sage” nuts (I think less than dozen in Missouri back then) and IPv6 troubles brought me to this group.

Yet another breach report

- Last week another ipv6-ops forum plea was made for firewalling on the internet.
- “Let me put it this way, I have personally found an anon-ftp server with company confidential documents on it, that was reachable from the outside without the owners knowledge, because there was a port-forward in the residential gateway that the owner wasn't actively aware of, and the NAS had anon-ftp turned on without the owners active knowledge.”

Continued

“So google had indexed all files on this NAS. I contacted the person (did some digging using pictures etc on this NAS) via their employer, and talked to the person who had no idea.

Now, with unfiltered IPv6 it would be harder to actually find this NAS, but once found, there is no need for port forward for it to be reachable from the Internet. That's why most ISPs have chosen to have stateful filtering toward the customers by default.”

Security Issues

Just take a look at many university networks. The ones I know use public IPv4 space, no NAT and many times not with firewalls. Now take one of those scanner / printer thinks with anon FTP saving all documents scanned on their local disk drive. Or power full laser with a power supply accessible via SNMP private. I think many people are accustomed to the "security" they get from NAT and don't think that there is anything else.

End

http://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/ip-multicast/whitepaper_c11-508498.html

http://ipv6friday.org/blog/2011/12/ipv6-multicast/

http://gobble.de.goop.maplepark.com/IPv6_Multicast.pdf (.odp or .ppt or .pptx also)

Thanks- Any Questions?