

UKERNA  
IP Multicast Mini Workshop  
Intra-domain Multicast

Hands-on Lab Exercises

Networkshop 2006

# Laboratory Overview

- Study the lab network topology
- Get familiar with Windows XP workstation platform and Cisco 2801 router platform
- Configure multicast on your edge and access router
- Configure a PIM Rendezvous Point (RP)
- Run ssm ping/asm ping tests
- Observe edge router multicast states

This set of exercises offers an overview of IP multicast in operation from a site administrator's point of view.

Our aim in setting the exercises is to take you through the multicast perspective from the bottom up, i.e. to begin with configuring IP multicast on a subnet including the on-link router.

You'll then configure the access router for your Team network, allowing multicast to flow between the access router, edge router and workstation systems, and configure the uplink on your edge router for multicast.

Then, you will define a PIM-SM Rendezvous Point (RP) for your domain on your access router, and configure your edge router to use the RP.

Having done that, you can use the ssm ping and asm ping to study how traffic flows within the domain, and how state changes in multicast routing occur on the routers.

# Getting help

- External network access
- Cisco IOS versions in use
  - 12.4T series
- Online reference
- Ask the helpers! 😊



Your workstations should have external IPv4 access; thus you can access remote web resources for help. Alternatively there are lab helpers to hand who should be able to assist you.

The lab routers are Ciscos running IOS12.4T.

Places to look when seeking on-line documentation include:

\* Cisco Software Release 12.4T Command References  
<http://www.cisco.com/univercd/cc/td/doc/product/software/ios124/>

\* Cisco Software Release 12.4T Debug Commands Reference  
[http://www.cisco.com/univercd/cc/td/doc/product/software/ios124/124cr/hdb\\_r/](http://www.cisco.com/univercd/cc/td/doc/product/software/ios124/124cr/hdb_r/)

Note that you're not expected to have to finish everything in this lab - it's designed for you to experiment at your own pace.

# Cisco IOS

- Knowledge of Cisco IOS would be useful. If new...
- Most configuration is undertaken with level-15 privileges (reached by "enable" on the console)
- Commands tab out, type "?" to show possibilities, a good starting point is "show ?", e.g. "show ip interfaces"
- Use "config terminal" to enter configuration mode.

The workgroup routers start off in a state in which you can access their CLI using telnet. The exec password you need is "nws\$password". There's only so much you can look at in IOS without having level-15 privileges, however it is handy to familiarise oneself with the environment by playing with a few commands - it's hard to break Ciscos without level-15 access!

"show ?" lists a set of possible command completions for the current context are displayed. Depending on your current privilege level, different options will avail themselves of you. Try "show ip interfaces" and press return, and have a look over the output.

The command that promotes an interactive session to level-15 privileged status is "enable". The enable password for the routers is "nws\$enable".

Once enabled, you can examine the running configuration, access much greater usage and debug information and, critically for this lab, configure the router.

To enter configuration mode, type "config terminal" (or just "conf t") and press return. Note how the prompt changes to reflect the new context - currently the router's global configuration context. Additional sub-contexts include interface definition mode where subsequent directives entered affect a particular interface; line definition mode to configure the administration and modem interfaces; etc. To step back up a level to the parent context (and from the global context back out of configuration mode and back into exec/interactive), type "exit" on a line of its own.

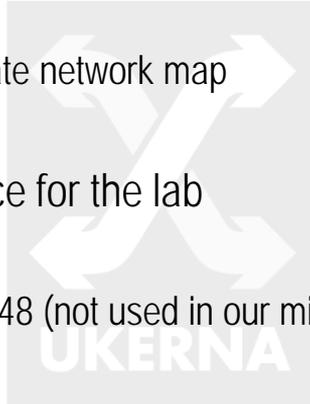
Once you've made configuration changes, the live running configuration can be dumped to console by typing "show running-config" in interactive mode. Likewise the configuration directives that will be applied post reboot is the 'startup-config'. To make the current live config persist (i.e. become the startup-config), enter "write memory" (or just "write") in interactive mode.

The laboratory helpers should have the base configurations on flash memory; so these can be reinstalled if you get in a knot.

A crash-course on IOS is beyond the scope of this workshop, however there are plenty of books and web articles that cover the elementary aspects of IOS configuration and use. As ever - just shout if you get stuck.

# Laboratory topology

- Topology
  - See the separate network map
  - Six teams, A-F
- IP address space for the lab
  - 193.61.85.0/24
  - 2001:630:23f::/48 (not used in our mini workshop)
- Admin privileges
  - See topology sheet



The laboratory is using a /24 size IPv4 prefix and a /48 size IPv6 prefix allocated by the JANET NOSC, resulting in every workstation and router having public and globally-routable IPv4 addresses. Since this allocation is not part of Hatfield's address space, the lab network is connected to the JANET multicast service via an IPv4 unicast tunnel from the core laboratory router (a Cisco 7206).

The laboratory is split into six Teams (A through F), with each group sharing a bench comprising two workstations, split into two subnets.

Each bench has a network topology map showing the initial state of the laboratory network.

You will have full administrator privileges on the workstations and level 15 access on your local workgroup router.

The windows administrator account is called "admin". The password is "qazwsx" (not particularly secure, admittedly, but this is a lab).

# Your Team cluster

- Each group has the following:
  - Cisco 2801 or 3825 access router
    - Two Ethernet ports (one uplink one downlink)
  - Cisco 2801 edge router
    - One Ethernet uplink
    - One single Ethernet subnet, one QuadFE subnet
  - Client nodes
    - Two Windows XP Service Pack 2 PCs, split into two subnets
  - Full IPv4 unicast connectivity

Please take a moment to study the laboratory network diagram.

You will see the six teams of people, Team A through Team F, each having a similar setup.

Each team has an edge router to which two host subnets are attached. In each case this is a Cisco 2801. Your two workstations are wired such that each workstation lies in one of the two subnets. The edge router represents a departmental router in a campus environment. Each subnet has a /29 (8 host addresses) so you have the capability to plug a laptop into one of the extra ports on the edge router with the quad Fast Ethernet card if you wish to do so (if you do, you can use the next IP address up from your Client2 node. There is no DHCP on the network.

Each team also has an access router connecting to their edge router downstream and a core Cisco 7206 upstream. Teams A-C have Cisco 2801 for this purpose while Teams D-F have a Cisco 3825. This router will be used as if it were a campus access router.

The 7206 is the core router, connecting downstream to each Team's access router and upstream to the JANET multicast service (via an IPv4 unicast tunnel). As the 'JANET-level' router, you will not have configuration access to this router.

The 7206 is running IOS 12.4(2)T4 with Advanced IP Services.

The 2801's and the 3825's are running IOS 12.4(4)T1 with Advanced IP Services.

Note that there is only one edge and one access router per group, so you will need to work together and rotate responsibility for configuring the router, etc.

# Current setup

- IPv4 unicast in 'provider'
  - Routing IPv4 unicast towards your access router
  - Routing IPv4 unicast from your access router to your edge router
  - Subnet routes (IPv4 unicast) for your workstation systems
- Stock IPv4 on your client subnets
  - Off-the-shelf Windows XP SP2
  - Manually configured IP settings (no DHCP)
- Next step: get the edge router ready
  - Configuring IP multicast

It is probably worth spending a little time getting familiar with the topology of the laboratory network.

Your first task will be to turn multicast on on your edge router, then to configure the access router, before running some multicast tests and observing router states.

# Multicast on the edge router

- Configure multicast on the edge router
- First inspect multicast setup
- Then turn on IP multicast
- Configure SSM default
- Disable IGMP snooping



First you should check on the edge router that no multicast is configured and be confident that no multicast routes are present. This should be the default setup/status when you site down to configure the edge router for the first time.

To do this, enter enable mode and use

```
show ip mroute
```

To turn on IP multicast, enter configure mode and then use the global command

```
ip multicast-routing
```

You should also tell IOS to treat the proper default range (232.0.0.0/8) as SSM

```
ip pim ssm default
```

And to avoid any potential issues with 'smart' Layer 2 snooping, you can disable Internet Group Management Protocol (IGMP) snooping

```
no ip igmp snooping
```

To leave config mode, type

```
exit
```

# Quad FE cards and IGMP snooping

- Edge router interfaces
  - One FE
  - One Quad FE
- A vlan has been configured spanning the Quad FE interface
  - IOS can be configured to do IGMP snooping to be 'smart' about which multicast traffic is flooded to which physical interfaces
  - For the sake of the workshop, we keep the setup 'clean'

The Cisco 2801 edge routers have a quad fast Ethernet card for one of the two workstation subnets; in principle IGMP snooping can prevent multicast flooding to switched Ethernet ports without interested hosts, but in practice use of IGMP (and MLD for IPv6) snooping needs to be considered carefully.

# Multicast on the host interfaces

- You now need to enter interface-specific commands to enable multicast on the edge router host subnets
- Single FE port:
  - Turn on PIM-SM
  - Turn on IGMPv3
- Quad FE port:
  - Turn on PIM-SM
  - Turn on IGMPv3

Having turned on multicast on the router, we now need to configure multicast on each workstation subnet.

First, we select the interface to apply the commands to; enter configure mode and then use

```
int fa0/1
```

will select the single FE interface. Then configure PIM SMI and configure IGMPv3 (so that SSM will work; with just IGMPv2 there will be no SSM capability)

```
ip pim sparse-mode  
ip igmp version 3
```

Type

```
exit
```

To then step out of the interface configuration level/mode.

Then apply the same configuration to the quad FE interface

```
int vlan1  
ip pim sparse-mode  
ip igmp version 3  
exit
```

You can at any time review the multicast routes that your edge router knows about with

```
show ip mroute
```

when outside of the config mode.

From any level you can use 'end' to exit the config mode.

# Edge Router uplink

- The next step is to enable multicast on the edge router uplink
  - Turn on PIM
- We don't need to enable IGMPv3, as that is a host-to-router protocol and this link is router-to-router
- We will then move on to configure your Team's access router.

We now enable multicast on the uplink.

Select the interface, and enable PIM:

```
int fa0/0
ip pim sparse-mode
```

There is no IGMPv3 on the router-to-router uplink.

# Multicast on access router

- There are three steps to configuring the access router
  - First, turn on multicast
  - Second, enable multicast interfaces (PIM)
  - Third, configure the RP
- Here we do the first two steps.
- Note Teams A-C have a Cisco 2801 access router, while Teams D-F have a Cisco 3825, so the interfaces will be slightly different
  - fa0/1 (on 2801) as opposed to gi0/1 (on 3825)

First, turn on multicast on the router at the conf level:

```
ip multicast-routing  
ip pim ssm default
```

Then enable multicast on the access router's internal domain interface:

For the 2810 (Teams A-C):

```
int fa0/1  
ip pim sparse-mode
```

For the 3825 (Teams D-F):

```
int gi0/1  
ip pim sparse-mode
```

The next step is to configure a PIM RP on the access router.

# Configuring the RP

- The RP function runs on one router in the PIM domain
  - We don't consider failover or Anycast-RP here
- Need to pick an RP address
- Next, we configure the RP

For multicast senders and receivers to meet in a campus scale PIM domain we need an RP configured. While this may seem overkill for our lab network, we couldn't expect Cisco to loan us up to 50 routers per Team ☺

The common thing to do these days, is to configure the RP address statically on all routers, and put that address on a loopback interface on the router chosen to be the RP. Then you can easily move the RP another router as needed. Also, when using Anycast-RP, that address might be on the loopback on multiple routers.

You typically also have a loopback address per router that is used for management (addresses on physical interfaces are only available if interface is up), often also used for BGP peerings, as ID in multiple routing protocols etc. It's nice to keep that loopback address separate from the RP one that might move or be in multiple places.

The RP address needs to be a unicast address within your IP allocation.

Your recommended Team RP address is listed on the network topology diagram in the Team section (though other addresses could be picked).

# Configuring the RP address

- First we need to create a loopback interface and assign it an IP address (that will be the RP address)
- We need to configure the RP address on the edge and access router
  - See the topology diagram for allocated RP addresses

We need to configure the RP address on both the edge and access router.

The access router will be the RP; we will first create a loopback interface to bind the RP's IP address to:

```
int loopback1
ip address <rpaddress> 255.255.255.255
```

While still on the access router, then configure the RP address:

```
ip pim rp-address <rp_address>
```

Finally, on the edge router:

```
ip pim rp-address <rp_address>
```

Remember, your RP address is on the topology diagram for your Team.

# IOS commands to try

- There are some specific IOS commands you can use to view the PIM state on the router(s)
  - Look at the PIM interfaces
  - Look at the PIM neighbors
  - Check RP information
  - Show how router is doing RPF
  - Look at IGMP group information
  - Look at the multicast routing table

Here are some commands you can try before, after or while running asmping.

Look at the PIM interfaces:

```
show ip pim interface
```

Show your PIM neighbours:

```
show ip pim neighbor
```

Look at the RP information:

```
show ip pim rp
```

For RPF information:

```
show ip rpf
```

To look at IGMP groups on the router:

```
show ip igmp groups
```

And in detail:

```
show ip igmp groups detail
```

To verify multicast is enabled

```
show ip multicast
```

To see multicast route information use

```
show ip mroute
```

To show counts on the routes

```
show ip mroute count
```

# About ssm ping

- Testing multicast connectivity
- Use ssm ping
  - Run daemon ssm pingd server on one system
  - Run client ssm ping on another
- Client signals to server that it wishes to receive multicast
  - Server responds with multicast (and unicast)

The ssm ping tool (and its integral *asm ping* counterpart) is very useful for testing multicast connectivity. By running these tools, one can send a message to a host on the Internet (which is running the *ssm ping/asm ping* server) and one will see whether one can receive (unicast and) multicast from the server back to the ssm ping client host.

The tool thus allows us to generate a multicast flow from one host (running the daemon) to another (the client running ssm ping).

The ssm ping client signals to the server on UDP port 4321.

The client will join (S,G) where S is its IP address and G is set by ssm ping to be 232.43.211.234.

The *ssm ping* package is pre-installed for you on the lab workstations, but is also available from <http://www.venaas.no/multicast/ssm ping/> for Windows and Unix platforms.

# Using ssm ping

- We now wish to observe multicast flows between the subnets attached to our edge router
  - We'll use ssm ping
- Running ssm ping between workstations
  - Client to server
  - Check multicast route states

The ssm ping package should be preinstalled on your workstations.

Choose one workstation to be the server and one to be the client.

Start ssm pingd on the server

```
ssmpingd.exe
```

Then run ssm ping from the client to the server

```
ssmping-0.8.1.exe <server_ipaddress>
```

You should see unicast and multicast replies.

Log in to the edge router and check the multicast routes

```
show ip mroute
```

What new route entries do you see? Can you identify them for specific (S,G) pairs?

Have a look at the host routing table on the XP workstation. At the command line use

```
route print
```

Stop the ssm ping client.

Check the edge router multicast route status.

# Exploring the PIM domain

- Having configured the RP, you can also test the PIM-SM multicast operation using asmping
  - asmping is the ASM variant of ssm ping
- Run asmping between two workstations
  - Observe PIM-SM protocol in action
  - Observe router multicast states/routing tables

While our PIM domain setup is relatively simple, it allows you to get an insight into the operation of PIM-SM as described in the preceding theory session.

We'll use asmping between a client and server on each of your Team's workstation subnets to generate multicast (\*,G) traffic.

While asmping is running, you can look at the multicast routing tables and related state on both routers to see how they change over time, and after asmping is stopped.

# Using asmping

- The asmping client works similarly to ssm ping
  - But you need to specify a group address to use
    - `asmping <multicast_group> <server_ip>`
  - An example group to use could be 224.1.2.234
- Note that the group address you use must end .234 (a 'security' feature of asmping).

Start ssm pingd on your chosen ssm ping server host:

```
ssmpingd.exe
```

Run asmping from one client workstation to the server workstation:

```
asmping-0.8.1.exe <group> <server_ip>
```

Study the router multicast routing information.

Stop the asmping flow.

Look at the subsequent router statuses.

## Aside: Ethereal

- Windows packet analyser
  - Preinstalled for you
- Allows capture of multicast traffic
- Useful to understand multicast flows on the network
- Allows filtering
  - e.g. all multicast traffic
  - Traffic to/from a specific host
- Have a go if time permits

In order to view the interesting multicast packets on the network, you could also use a packet analyser.

For the exercises, and the time allocated, we felt that using Ethereal may cause time pressure. However, if you've got this far before the session has ended, feel free to rerun `asming` or `ssmping` and do Ethereal packet dumps and analysis.

We have pre-installed Ethereal on the Team workstations.

When running Ethereal you will by default see all traffic. In our lab that will be limited; in a real environment it will be much busier.

Ethereal allows you to filter specific traffic to view. You can click on 'Capture' on the menu bar, then select Capture Filters, which pops up a window in which you can select the "New" button, and enter filter strings for filtering. You can experiment with these.

As a filter string

`ip multicast`

Should capture all multicast traffic.

To capture all traffic to or from a host, use

`host <ipaddress>`

(though note each workstation is in its own subnet in our lab, unless extra laptop devices are plugged in to the QuadFE ports)

# Lab Overview

- What did we do
- What did we learn
- What would be different in the real world

