

# JANET

## IPv6 Hands-on Workshop

Lab 3: **IS-IS Routing**

UKERNA, Lancaster University  
and University of Southampton, 2006

# Laboratory Overview

- Enable IS-IS on your Edge Router
- Set Level 2 IS-IS Transport Mode
- Apply route aggregation for your /60 prefixes
  - We have a /56 for the workshop (2001:630:81:0400::/56)
  - Each group has a /60 (16 subnets)
  - Each group subnet has a /64
- View discovered routes to other teams branches

# IS-IS

- For this exercise we will all work together, like we are all in the same admin domain
  - e.g. a university campus
- We will then configure IS-IS between our routers
  - We only configure one common backbone area as this is sufficient unless there is a huge number of routers
- We assume we all are in Autonomous System (AS) 65000 (**fde8** in hex) and we use **0** for the area id.
- On the central router we have an IPv4 loopback address of **148.88.147.220**
  - Thus our IS-IS id is 47.fde8.0000.1480.8814.7220.00

September 2005

UKERNA IPv6 Workshop

3

Routing tables for IPv6 have been removed from the head router and you should find that attempting to communicate over IPv6 to the head router and beyond will not currently be working (e.g. ping6 2001:630:81:4X0::1). The static IPv6 route on your edge router configured yesterday has also been removed.

The head router is now supporting IS-IS for routing and has been configured with the IS-IS id of 47.fde8.0000.1480.8814.7220.00. We are going to work together in this exercise and pretend we are in the same admin domain.

NOTE: IS-IS ids can be of arbitrary length thus in an IPv6-only deployment you could use a full IPv6 address if you chose to do so.

## Step1 - Global Configuration

- Configure the IS-IS id on your edge router
- You need to specify **net** with an NSAP address as follows
  - It will start with **47.fde8.0000**
  - The only difference from the address on the previous slide will be the IPv4 address used
  - You can pick any of your IPv4 addresses, we suggest using the address of the upstream interface (**FastEthernet-1/0/0**)
  - normally a loopback address would be a good choice, but no loopback has been configured here
  - You can make use of an IPv6 address in the IS-IS id, remember to include all 0's
- Set it to use level 2 transport only

September 2005

UKERNA IPv6 Workshop

4

### CISCO

The first stage is to enable IS-IS routing by entering the command:

```
router isis
```

Then take your net with NSAP address and specify this by executing the following:

```
net 47.fde8.0000.1480.8814.7254.00
```

To specify the level of transport to be level 2 only:

```
is-type level-2-only
```

### JUNIPER

In JUNOS the router uses its lo0 interface for application of global settings, we have met this before when applying firewall filters in the previous lab. Thus in the same way the NSAP identifier is applied to this interface:

```
from [edit interfaces]
  set lo0 unit 0 family iso address 47.fde8.0000.1480.8814.7254.00
```

The IS-IS type can be defined per interface, or globally in [edit protocols isis] as is the case here. By entering [edit protocols isis] we enable IS-IS on your router and can specify the type with the following command:

```
from [edit protocols isis]
  set level 1 disable
```

### BOTH

Your router is now globally configured however no changes will be noticeable until we enable the interfaces, this is the next step.

## Step2 - Enabling the Up Interface

- You need to enable all the interfaces which you wish to be advertised using IS-IS
  - In this case just your up interface
  - Never an interface facing a host
- Then we'll turn off IPv4 IS-IS routing (next page) before committing the change

September 2005

UKERNA IPv6 Workshop

5

### CISCO

Enabling IS-IS in IOS is done per interface and per address group, so to enable IPv4 IS-IS routing we enable "ip router isis" on each interface. To enable IS-IS routing, for each interface enable the following:

```
ip router isis
ipv6 router isis
```

### JUNIPER

To enable IS-IS, interfaces first need to support the iso family, just like was done previously for the lo0 interface. Use the following command, assuming your up interface is fe-0/0/0:

```
from [edit interfaces]
  set fe-0/0/0 unit 0 family iso
```

Once the interfaces are configured we finally need to add all these to the IS-IS global setting.

```
from [edit protocols isis]
  set interface lo0
  set interface fe-0/0/0
```

**DO NOT COMMIT** the configuration at this stage... we will turn off IS-IS IPv4 routing first - next page!

## Step2 - Enabling the Up Interface (ctd)

- Disable ISIS from doing IPv4 unicast routing
  - Static routes already there.
- Disable IPv4 ISIS on the interfaces
- Check to see if you are connected!

September 2005

UKERNA IPv6 Workshop

6

IPv6 and IPv4 routing options in JUNOS are also set from within [edit protocols isis] and can be global or applied per interface.

Global:

```
from [edit protocols isis]
  set no-ipv4-routing
```

Per Interface:

```
from [edit protocols isis]
  set interface fe-0/0/1 no-ipv4-unicast
```

Note the slight different in options between the global and per interface. Finally to see if you are connected try the following command:

```
show isis adjacency
```

### **CISCO**

To see if you are connected at this point on IOS try the following command:

```
show isis neighbors
```

## Verifying that IS-IS works

- Check that you have an IS-IS neighbour
- For checking the database try
  - `show isis database`
  - `show isis database detail`
- With the latter you should see
  - The IS-IS state for your router including the IPv4 and IPv6 information for your upstream interface, including any prefixes you are advertising.
  - State from core router with all its downstream interfaces
  - Information from some of the other group routers
    - Unless you are the first. In this case wait until you can see some of the others
- Show your routing table to verify that you now have learned some additional L2 routes.

September 2005

UKERNA IPv6 Workshop

7

At this stage we could skip all of this section and try the ping test mentioned on the first lab slide to see if your connectivity has returned. However there are also several commands which we will need later when we look at aggregation.

As a repeat of the last point on the previous slide we can firstly check that we have some IS-IS neighbours. Further information can also be obtained by looking at the ISIS database.

### **BOTH**

```
show isis database
show isis database detail
```

We can also view any routes you have learned at this point and verify that these are L2 routes.

### **CISCO**

```
show ipv6 route
```

### **JUNIPER**

```
show route table inet6
show isis route
```

We should see both IPv4 and IPv6 routes being advertised by IS-IS here, however we have IPv4 routed statically. Under IOS we can just remove the "ip router isis" command from the interface configs. However under JUNOS, this is done using a policy (see next slide).

## Step3 - Including routes for non IS-IS interfaces

- You should have now learnt some new routes from the head router. For IS-IS to be useful, we need to also announce routes which are available through your edge router.
- We will do this only for IPv6.
- To achieve this we need to tell the router to redistribute its connected hosts.
- At the end of this your router will be advertising 3 /64 prefixes.
  - Both the upstream and the two downstream

September 2005

UKERNA IPv6 Workshop

8

### CISCO

To do this we first enter  
router isis

Inside there do  
address-family ipv6  
redistribute connected

If you have static routes for external networks not using IS-IS you could have used “redistribute static” to include those.

### JUNIPER

You may have noticed Juniper like using policies to define routing and firewall options and unfortunately this operation is no exception. Firstly we need to create a policy which will only redistribute the connected IPv6 hosts.

```
From [edit policy-options policy-statement redistribute-connected]
  set from protocol direct
  set from family inet6
  set then accept"
```

Now returning to the IS-IS global options we apply this policy:

```
From [edit protocols isis]
  set export redistribute-connected
```

## Step4 - Aggregation

- For the groups a-f (or 1-6) the prefixes for the two downstream interfaces are **2001:630:81:4X2::/64** and **2001:630:D0:4X4::/64** where **X** is the group number (1-6)
  - Actually, each group has a **2001:630:D0:04X0::/60** at their disposal
- We will now perform this aggregation so that each group only announces their /60 prefix

September 2005

UKERNA IPv6 Workshop

9

### CISCO

To do this enter,

```
router isis
address-family ipv6
```

Inside there do:

```
summary-prefix 2001:630:D0:4X0::/60
```

(where X is the group number)

By repeating "show isis database detail" you should see that you now announce this /60 instead of your downstream /64s. Repeat this until you also see that the /64s from other groups are replaced with /60s. Static routes etc. can also be aggregated this way.

### JUNIPER

Like IOS we first need to add the aggregate route we wish to apply:

```
From [edit routing-options rib inet6.0]
set aggregate route 2001:630:81:4X0::/60
```

Then to cause this to take effect through IS-IS we need to make a new policy:

```
From [edit policy-options policy-statement v6-aggregation]
set term one from protocol aggregate
set term one then accept
```

## Step4 – Aggregation (ctd)

- Finish applying the aggregation so that each group only announces their /60 prefix

September 2005

UKERNA IPv6 Workshop

10

Applying the changes at this state by activating the policy would advertise a new aggregated route however the /64's will still be advertised, to prevent this we extend our policy with a second term.

```
From [edit policy-options policy-statement v6-aggregation]
  set term two from route-filter longer 2001:630:81:4X0::/60
  set term two then reject
```

To activate this policy we need to add it to the global IS-IS options as before:

```
From [edit protocols isis]
  set export v6-aggregation
```

This appends the v6-aggregation policy **after** the redistribute-connected policy has already run. This means that you will still be advertising the three /64 prefixes (see "show isis database detail") as well as the new /60. To prevent this the aggregation policy need to be run first. To ensure this happens execute the following:

```
From [edit protocols isis]
  insert export-aggregation before redistribute-connected
```

Try viewing the database again to ensure you are now only advertising the /60 prefix.

### **BOTH**

Viewing your IS-IS statistics at this stage should show these latest set of changes and reveal if you were first to complete.

## Summary

- IS-IS Configuration and Operation
- Redistribution of Connected Hosts/Networks
- Aggregation
  
- Next: BGP and Multicast