

# IS-IS Lab

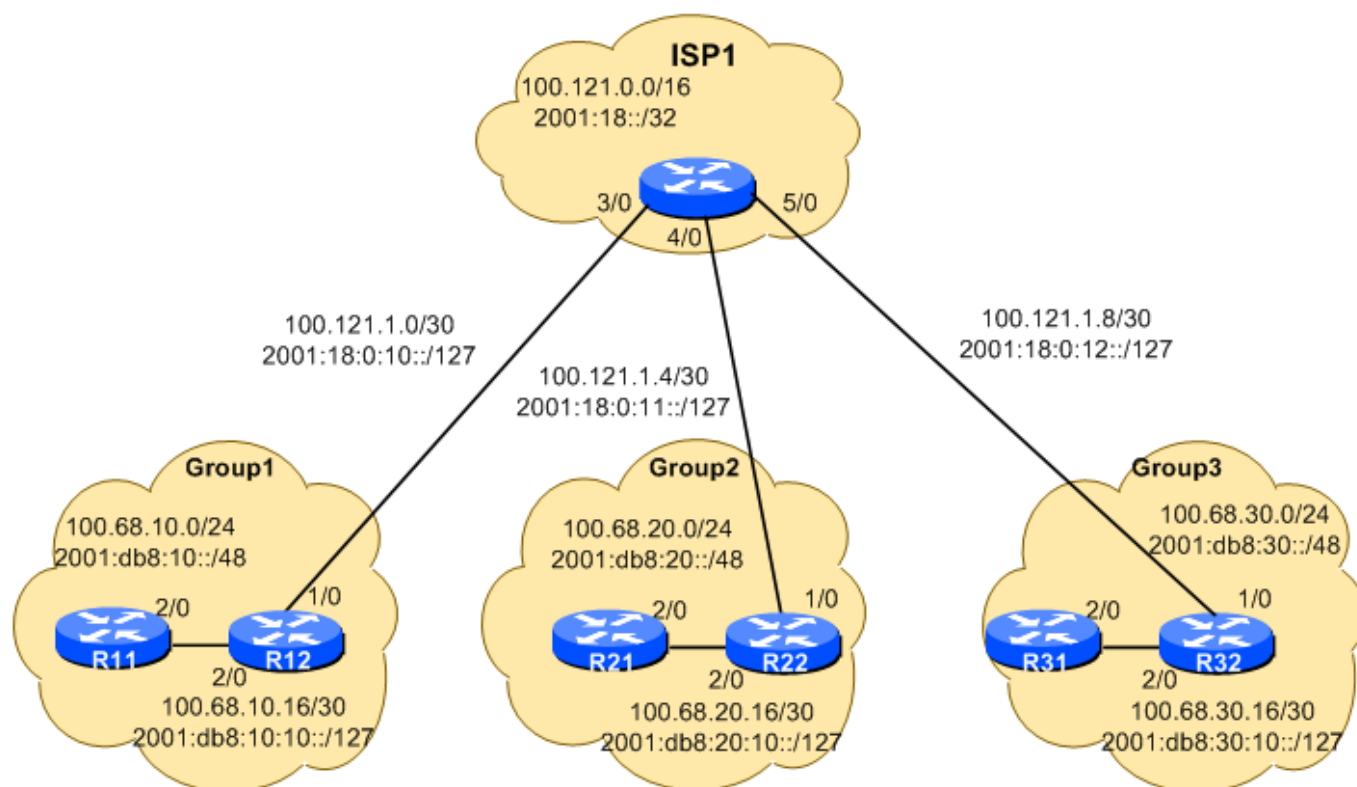
## Introduction

The purpose of this exercise is to:

- Enable IS-IS to exchange internal routing information
- Configure static routing towards a service provider


### Topology of Module 1

As a reminder, here is the topology of Module one. It is made up of three Groups (1,2,3), and their ISP. As we go through the workshop, we will add an NREN to this Module as well.



### Topology of Module 2


As a reminder, here is the topology of Module two. It is made up of three Groups (4,5,6), and their ISP. As we go through the workshop, we will add an NREN to this Module as well.



## Topology of the whole Lab

The entire workshop lab is interconnected as in the following diagram - the two ISPs peer with each other at the Internet Exchange Point (IXP).

The network configuration is designed to be modular to allow the lab to grow as needed depending on the number of participants. Each module contains 1 ISP and 3 customer networks (universities, etc). Modules are be interconnected as we will see.



# Creating Loopback and Null Interfaces

## Loopback interface

Create the Loopback interface. We will need it later on for BGP (and indirectly for IS-IS as well).

**On RX1:**

```
interface loopback 0
 ip address 100.68.X0.1 255.255.255.255
 ipv6 address 2001:db8:X0::1/128
```

do the same for RX2 (obviously, using different addresses).

## Null Interface

Create Null interface and disable unreachable processing. This is a best practice configuration which ensures that the router will not send ICMP unreachable messages for traffic which lands on the Null interface.

```
interface null 0
 no ip unreachable
```

## Introducing Dynamic Routing

Try pinging the loopback addresses of your neighbour

```
RX1# ping 100.68.X0.2      <- RX2 loopback
RX1# ping 2001:db8:X0::2   <- RX2 loopback
```

What happens?

Can you explain why?

## Set up IS-IS

Now we will configure a new IS-IS routing process.

We will use the string "asX0" as the IS-IS process identifier for routers RX1 and RX2. This identifier is local to the router, so it doesn't need to match the process identifier of a neighbouring router. However, it is strongly recommended that you use the same number throughout your network. (For example, Group One would use the string "as10". Notice that this number looks the same as the AS number for your group - but there really is no relationship between them.)

IS-IS requires an NET to uniquely identify each router in the network. The NET needs to be configured

manually, and common convention is to use embed the router's loopback address. The NET should be 49.0001.xxxx.00 where xxxx is the router loopback IPv4 address. For example, for R11 the loopback is 100.68.10.1, and for R12 the loopback is 100.68.10.2. So the NETs for these two routers will be 49.0001.1000.6801.0001.00 for R11 and 49.0001.1000.6801.0002.00 for R12.

*Hint: A nice trick for converting the loopback interface address into the NET address is to take the loopback address and put the missing leading zeroes in. For example, R11 loopback address is 100.68.10.1 and becomes 100.068.010.001 putting in the missing zeroes. Then rather than having the dot after every third character, move it to be after every fourth character. So 100.068.010.001 becomes 1000.6801.0001.*

Finally, we will run our entire network as L2 only - there is no requirement for L1 until the network becomes really large (>400 routers).

### On RX1:

```
router isis asX0
 net 49.0001.1000.680X.0001.00
 is-type level-2-only
```

## Wide Metrics

We also set the metric-style to wide. IS-IS supports two types of metric, narrow (historic now and not suitable for modern networks) and wide. Cisco IOS still defaults to narrow metrics, so we need to enter configuration to change this to wide.

```
router isis asX0
 metric-style wide
```

## Default Metrics

The default metric in IS-IS is 10, on all interfaces irrespective of their physical bandwidth. It is considered best practice these days to change the default metric from 10 to a very high value, for example 100000, so that there are no network outages caused by misconfiguration of newly introduced routers, or misconfigured interfaces which could accidentally and unintentionally take full traffic load.

```
router isis asX0
 metric 100000
```

## Multi-Topology

We also need to activate multi-topology IS-IS so that we can support IPv6 across our network. Multi-topology allows IPv6 to be deployed incrementally, but it also reduces the chances of network outages or accidents if IPv6 configuration is omitted on any backbone links.

```
router isis asX0
```

```
address-family ipv6
multi-topology
```

## Overload

We also need to let IS-IS know about when a router has newly rebooted, that it should not be considered for carrying traffic until iBGP (in the next lab) is up and running. (With no iBGP, as the case is now, the configuration below has no effect.)

```
router isis asX0
set-overload-bit on-startup wait-for-bgp
address-family ipv6
set-overload-bit on-startup wait-for-bgp
```

*The overload bit is used to signal to IS-IS neighbours that the router is “overloaded” and should not be used for traffic. However, in today's networks, network operators use the overload bit to indicate that the router should not carry traffic until BGP is operational.*

## Neighbour Authentication

Neighbour authentication is highly recommended. First we create a suitable key-chain:

```
key chain isis-asX0
key 1
key-string nsrc-isis
```

And then we apply the key-chain to the IS-IS process (which means that neighbour authentication will apply on all interfaces where adjacencies will be established):

```
router isis asX0
authentication mode md5 level-2
authentication key-chain isis-asX0 level-2
```

## IS-IS on Interfaces

Now configure IS-IS on the interfaces *where adjacencies need to be established*. Note that point-to-point connections using broadcast media such as Ethernet should be declared as point-to-point.

### On RX1 and RX2:

```
interface GigabitEthernet2/0
isis network point-to-point
ip router isis asX0
isis metric 2
ipv6 router isis asX0
isis ipv6 metric 2
```

## Announcing Loopback addresses

Finally, we need to announce the loopback interface address by IS-IS as well:

```
router isis asX0
passive-interface Loopback0
```

## CHECKPOINT

Now try the following show commands:

```
sh clns neighbor      : show CLNS adjacencies
sh isis neighbor      : show IS-IS adjacencies
sh ip route           : show the IPv4 routes in routing table
sh isis ip rib        : show the IPv4 ISIS RIB
sh ipv6 route         : show the IPv6 routes in routing table
sh isis ipv6 rib      : show the IPv6 ISIS RIB
```

Repeat the last ping tests. Can you ping the loopback address of the neighbouring router now?

## Static default routes

In the previous exercise we configured a large number of static routes on RX1 and RX2 so that we could reach the other groups in the lab. Clearly, the more groups we add, the harder it is going to be to maintain this whole static route structure. Which is why we do not use static routes very much in the day to day operational Internet.

## Removing explicit Statics

Search through your configuration and look for the static routes that were created in the previous lab:

```
RX1# show run | include ^ip route
```

will show you all lines beginning with *ip route*, the static routes set up earlier. There is a similar command you can use for IPv6 - can you guess what it is?

To remove the routes, simply put the *no* keyword in front of the static route commands you used earlier - or use a text editor and take the output from the above command. For example, removing the two static routes for ISP1 and ISP2 address space on R11 is achieved like this:

```
no ip route 100.121.0.0 255.255.0.0 100.68.10.18
no ip route 100.122.0.0 255.255.0.0 100.68.10.18
```

**Do the same for all the other static routes configured on RX1 and RX2.**

## Setting up Defaults

Configure static **default** routes to reach the outside world.

### On R22:

```
ip route 0.0.0.0 0.0.0.0 100.121.1.5
ipv6 route ::/0 2001:18:0:11::0
```

Also on RX2 we are going to originate a default route using IS-IS:

```
router isis asX0
default-information originate
!
address-family ipv6
 default-information originate
!
```

This will originate a default route into IS-IS so long as there is a default route in the routers global RIB. (In other words, so long as the next hop of the default route we just configured on RX2 is reachable, IS-IS will originate a default route towards its neighbours.)

### Q. Why do we not need to configure any static routes on RX1?

## Testing Routing

Do some ping and traceroute tests. For example, on R11, try pinging Group 2 and Group 3 routers.

```
R11# ping 100.68.20.1
R11# ping 100.68.30.1
R11# traceroute 100.68.20.1
R11# traceroute 100.68.30.1
```

Can you reach the routers in other networks?

Don't forget to save your configurations before completing this lab.

From:

<https://wiki.lpnz.org/> - **Workshops**

Permanent link:

<https://wiki.lpnz.org/doku.php?id=2015:ubuntunet-bgp:1-lab-isis>

Last update: **2015/11/17 14:25**

