Basic Topology and ISIS

- Objective: Create a basic physical lab interconnection with one ISIS Area. Ensure that all
- routers, interfaces, cables and connections are working properly.
- Prerequisites: Knowledge of Cisco router CLI, previous hands on experience.

The following will be the common topology used for the labs.

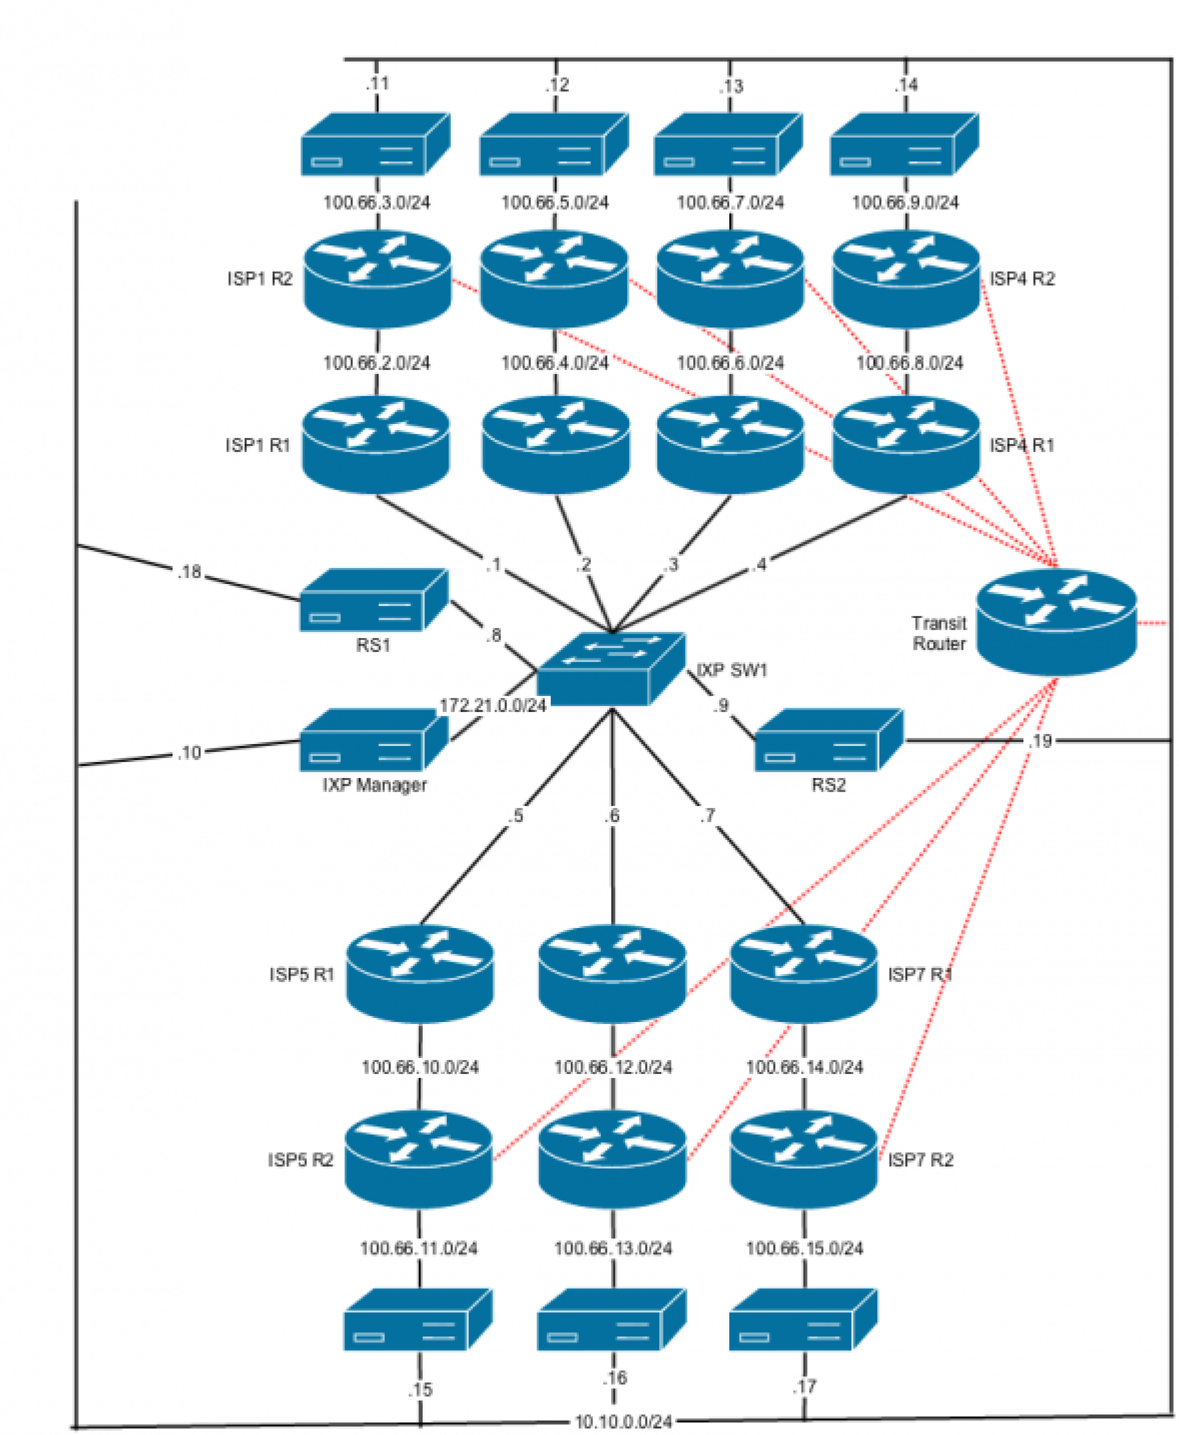

Figure 1 – ISP Lab Basic Configuration

Lab Notes

This workshop is intended to be run on a Dynamips server with the appropriate lab topologies set up. The routers in the Dynamips environment are using service provider IOS. The configurations and configuration principles discussed below will work on all Cisco IOS Release 12.4 onwards. Earlier Cisco IOS releases are not supported but will mostly work using the notes below; they will miss some of the features covered.

The purpose of this module is to construct the workshop lab and introduce everyone to the basic principles of constructing and configuring a network. An important point to remember, and one that will be emphasised time and again through out this workshop, is that there is a distinct sequence to building an operational network:

- After the physical design is established, the connections between the hardware should be built and verified.
- Next, the routers should have the base configuration installed, and basic but sufficient security should be set up.
- Next the basic IP connectivity be tested and proven. This means assigning IP addresses on all links which are to be used, and testing the links to the neighbouring devices.
- Only once one router can see its neighbour does it make sense to start configuring routing protocols. And start with the IGP (ISIS is chosen for this workshop). There is no purpose to building BGP while the chosen IGP (in this case ISIS) is not functioning properly. BGP relies on ISIS to find its neighbours and next hops, and an improperly or non-functioning ISIS will result in much time wasted attempting to debug routing problems.
- Once the IGP is functioning properly, the BGP configuration can be started, first internal BGP, then external BGP.
- Remember to RTFM. What is RTFM? It is critical that ISP Network Engineers fully utilise all information resources. The #1 source is the documentation. Read The F#$% Manual (RTFM) is the traditional phase used to inform engineers that the answer is in the documentation and go read it. We will use RTFM through out these exercises to highlight areas where the student should use the documentation for further deepening. There will be many new commands. Please refer to the IOS Command Reference on Cisco Connection Online (CCO.cisco.com) for details on each of these commands.
- Finally, documentation. Documentation is often overlooked or forgotten. It is an ongoing process in this workshop. If the instructor asks you to document something, either on the whiteboard in the class, or at the back of this booklet, it is in your best interests to do so. There can never be too much documentation, and documentation at the time of network design and construction canusually saves much frustration at a future date or event.

Lab Exercise

1. Routers and the Workshops participants.

This workshop is laid out such that a group of students will operate a pair of routers. Because we have a large group we have setup two separate IXPs, IXP1 and IXP2. There are 7 pairs of routers, attached to each IXP, where each pair represents an ISP. The Workshop Instructors will divide the routers amongst the workshop participants. Take note of which ISP and IXP combintaion you have been allocated to. In the following notes, a “router team” refers to the group assigned to one particular ISP.

2. Introducing the lab.

This workshop uses Cisco IOS routers running IOS, but on the Dynamips systems – Dynamips translates the Cisco 7200 router PowerPC processor instructions in IOS to those of the host system, allowing Cisco IOS images, and therefore network configurations, to be run on a host PC system (usual Linux or MacOS based). The lab will have been preconfigured by the instructors, allowing participants to enter the following exercises directly. Please read the following steps carefully.

3. Accessing the lab.

The instructors will assign routers to each class group, and will indicate the method of access to the Dynamips server. This will usually be by wireless – if this is the case, make a note of the SSID and any password required. Also make a note of the IP address (IPv4, as Dynamips only supports IPv4 access) of the Dynamips server.

Access to Dynamips will be by telnet, to a high port, which the instructor will specify. Each participant should ensure that their device has a suitable telnet client. Linux and MacOS system have access to a shell command prompt (or Terminal) programme, which allows telnet at the command line. Windows users can use the Windows “Command Prompt” with the telnet client there, but it’s notoriously unreliable. Better to install software such as Putty, TeraTerm, HyperTerm or similar third party telnet client.

Router	Connection details		Router	Connection details
r1-isp1-ixp1	telnet s1.ws.nsrc.org 2111		r1-isp1-ixp2	telnet s1.ws.nsrc.org 2121
r2-isp1-ixp1	telnet s1.ws.nsrc.org 2211		r2-isp1-ixp2	telnet s1.ws.nsrc.org 2221
r1-isp2-ixp1	telnet s1.ws.nsrc.org 2112		r1-isp2-ixp2	telnet s1.ws.nsrc.org 2122
r2-isp2-ixp1	telnet s1.ws.nsrc.org 2212		r2-isp2-ixp2	telnet s1.ws.nsrc.org 2222
r1-isp3-ixp1	telnet s1.ws.nsrc.org 2113		r1-isp3-ixp2	telnet s1.ws.nsrc.org 2123
r2-isp3-ixp1	telnet s1.ws.nsrc.org 2213		r2-isp3-ixp2	telnet s1.ws.nsrc.org 2223
r1-isp4-ixp1	telnet s1.ws.nsrc.org 2114		r1-isp4-ixp2	telnet s1.ws.nsrc.org 2124
r2-isp4-ixp1	telnet s1.ws.nsrc.org 2214		r2-isp4-ixp2	telnet s1.ws.nsrc.org 2224
r1-isp5-ixp1	telnet s1.ws.nsrc.org 2115		r1-isp5-ixp2	telnet s1.ws.nsrc.org 2125
r2-isp5-ixp1	telnet s1.ws.nsrc.org 2215		r2-isp5-ixp2	telnet s1.ws.nsrc.org 2225
r1-isp6-ixp1	telnet s1.ws.nsrc.org 2116		r1-isp6-ixp2	telnet s1.ws.nsrc.org 2126
r2-isp6-ixp1	telnet s1.ws.nsrc.org 2216		r2-isp6-ixp2	telnet s1.ws.nsrc.org 2226
r1-isp7-ixp1	telnet s1.ws.nsrc.org 2117		r1-isp7-ixp2	telnet s1.ws.nsrc.org 2127
r2-isp7-ixp1	telnet s1.ws.nsrc.org 2217		r2-isp7-ixp2	telnet s1.ws.nsrc.org 2227

Using the client, connect to the router you have been assigned; for example, to connect to the console port of r1-isp1-ixp1:

```
telnet s1.ws.nsrc.org 2111
```

or to r2-isp3-ixp2:

```
telnet s1.ws.nsrc.org 2223
```

Once connected, you will see the Dynamips response, followed by the login or command prompt of the router:

```
bash-3.2$ telnet s1.ws.nsrc.org 2111

Trying 10.10.0.241...
Connected to s1.ws.nsrc.org.
Escape character is '^]'.
Connected to Dynamips VM "r1-isp1-ixp1" (ID 2, type c7200) - Console port
Press ENTER to get the prompt.

....

User Access Verification

Username:
```

If the “Connected to Dynamips VM” won’t appear, even after hitting the Return key several times, please request help from the workshop instructors.

4. Basic Router Configuration

Each router will be named according to the table above: r1-isp1-ixp1, r2-isp1-ixp1, r1-isp2-ixp1, etc.

See r1-ixp6-isp1 config and r2-isp6-ixp1 config for details. Check that your routers have a similar config.

Hostname

Your routers have been given a basic configuration as follows:

```
Router> enable
Router# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# hostname r1-isp1-ixp1
Router1(config)#
```

Turn Off Domain Name Lookups

Cisco routers will always try to look up the DNS for any name or address specified in the command line. You can see this when doing a trace on a router with no DNS server or a DNS server with no inaddr.arpa entries for the IP addresses. We will turn this lookup off for the labs for the time being to speed up traceroutes.

```
Router1 (config)# no ip domain-lookup
```

Disable Command-line Name Resolution

The router by default attempts to use the various transports it supports to resolve the commands entered into the command line during normal and configuration modes. If the commands entered are not part of Cisco IOS, the router will attempt to use its other supported transports to interpret the meaning of the name. For example, if the command entered is an IP address, the router will automatically try to connect to that remote destination. This feature is undesirable on an ISP router as it means that typographical errors can result in connections being attempted to remote systems, or time outs while the router tries to use the DNS to translate the name, and so on.

```
Router1 (config)# line con 0
Router1 (config-line)# transport preferred none
Router1 (config-line)# line vty 0 4
Router1 (config-line)# transport preferred none
```

Disable Source Routing

Unless you really believe there is a need for it, source routing should be disabled. This option, enabled by default, allows the router to process packets with source routing header options. This feature is a well-known security risk as it allows remote sites to send packets with different source address through the network (this was useful for troubleshooting networks from different locations on the Internet, but in recent years has been widely abused for miscreant activities on the Internet).

```
Router1 (config)# no ip source-route
```

Usernames and Passwords

All router usernames should be isplab and all passwords should be lab-PW. Please do not change the username or password to anything else, or leave the password unconfigured (access to vty ports is not possible if no password is set). It is essential for a smooth operating lab that all participants have access to all routers.

```
Router1 (config)# username isplab secret lab-PW
Router1 (config)# enable secret lab-PW
Router1 (config)# service password-encryption
```

The service password-encryption directive tells the router to encrypt all passwords stored in the router’s configuration (apart from enable secret which is already encrypted). Note A: There is the temptation to simply have a username of cisco and password of cisco as a lazy solution to the username/password problem. Under no circumstances must any service provider operator ever use easily guessable passwords as these on their live operational network.

IMPORTANT: This sentence cannot be emphasized enough. It is quite common for attackers to gain access to networks simply because operators have used familiar or easily guessed passwords.

Note B: for IOS releases prior to 12.3, the username/secret pair is not available, and operators will have to configure username/password instead. The latter format uses type-7 encryption, whereas the former is the more secure md5 based encryption.

Enabling login access for other machines

In order to let you telnet into your router in future modules of this workshop, you need to configure a password for all virtual terminal lines.

```
Router1 (config)# aaa new-model
Router1 (config)# aaa authentication login default local
Router1 (config)# aaa authentication enable default enable
```

This series of commands tells the router to look locally for standard user login (the username password pair set earlier), and to the locally configured enable secret for the enable login. By default, login will be enabled on all vtys for other teams to gain access.

Configure system logging

A vital part of any Internet operational system is to record logs. The router by default will display system logs on the router console. However, this is undesirable for Internet operational routers, as the console is a 9600 baud connection, and can place a high processor interrupt load at the time of busy traffic on the network. However, the router logs can also be recorded into a buffer on the router – this takes no interrupt load and it also enables to operator to check the history of what events happened on the router. In a future module, the lab will configuration the router to send the log messages to a SYSLOG server.

```
Router1 (config)# no logging console
Router1 (config)# logging buffer 8192 debug
```

which disables console logs and instead records all logs in a 8192 byte buffer set aside on the router. To see the contents of this internal logging buffer at any time, the command “sh log” should be used at the command prompt.

Save the Configuration.

With the basic configuration in place, save the configuration. To do this, exit from enable mode by typing “end” or “<ctrl> Z”, and at the command prompt enter “write memory”.

```
Router1(config)#^Z
Router1# write memory
```

```
Building configuration...
[OK]
Router1#
```

It is highly recommended that the configuration is saved quite frequently to NVRAM. If the configuration is not saved to NVRAM, any changes made to the running configuration will be lost after a power cycle or virtual machine failure

Log off the router by typing exit, and then log back in again. Notice how the login sequence has changed, prompting for a “username” and “password” from the user. Note that at each checkpoint in the workshop, you should save the configuration to memory – remember that powering the router off will result in it reverting to the last saved configuration in NVRAM.

5. IP Addresses

This Module will introduce the basic concepts of putting together a sensible addressing plan for an ISP backbone. We are building one autonomous system out of the 14 routers we have in the lab. The RIRs are typically handing out IPv4 address space in /20 chunks (depends on which RIR region). For the purposes of this lab that our ISP has received a /23 of IPv4 space and a /60 of IPv6. Rather than using public address space, we are going to use a portion of 100.64/10 (RFC 6598 - carrier-grade NAT) for this lab. In the real world Internet, we would use public address space for our network infrastructure.

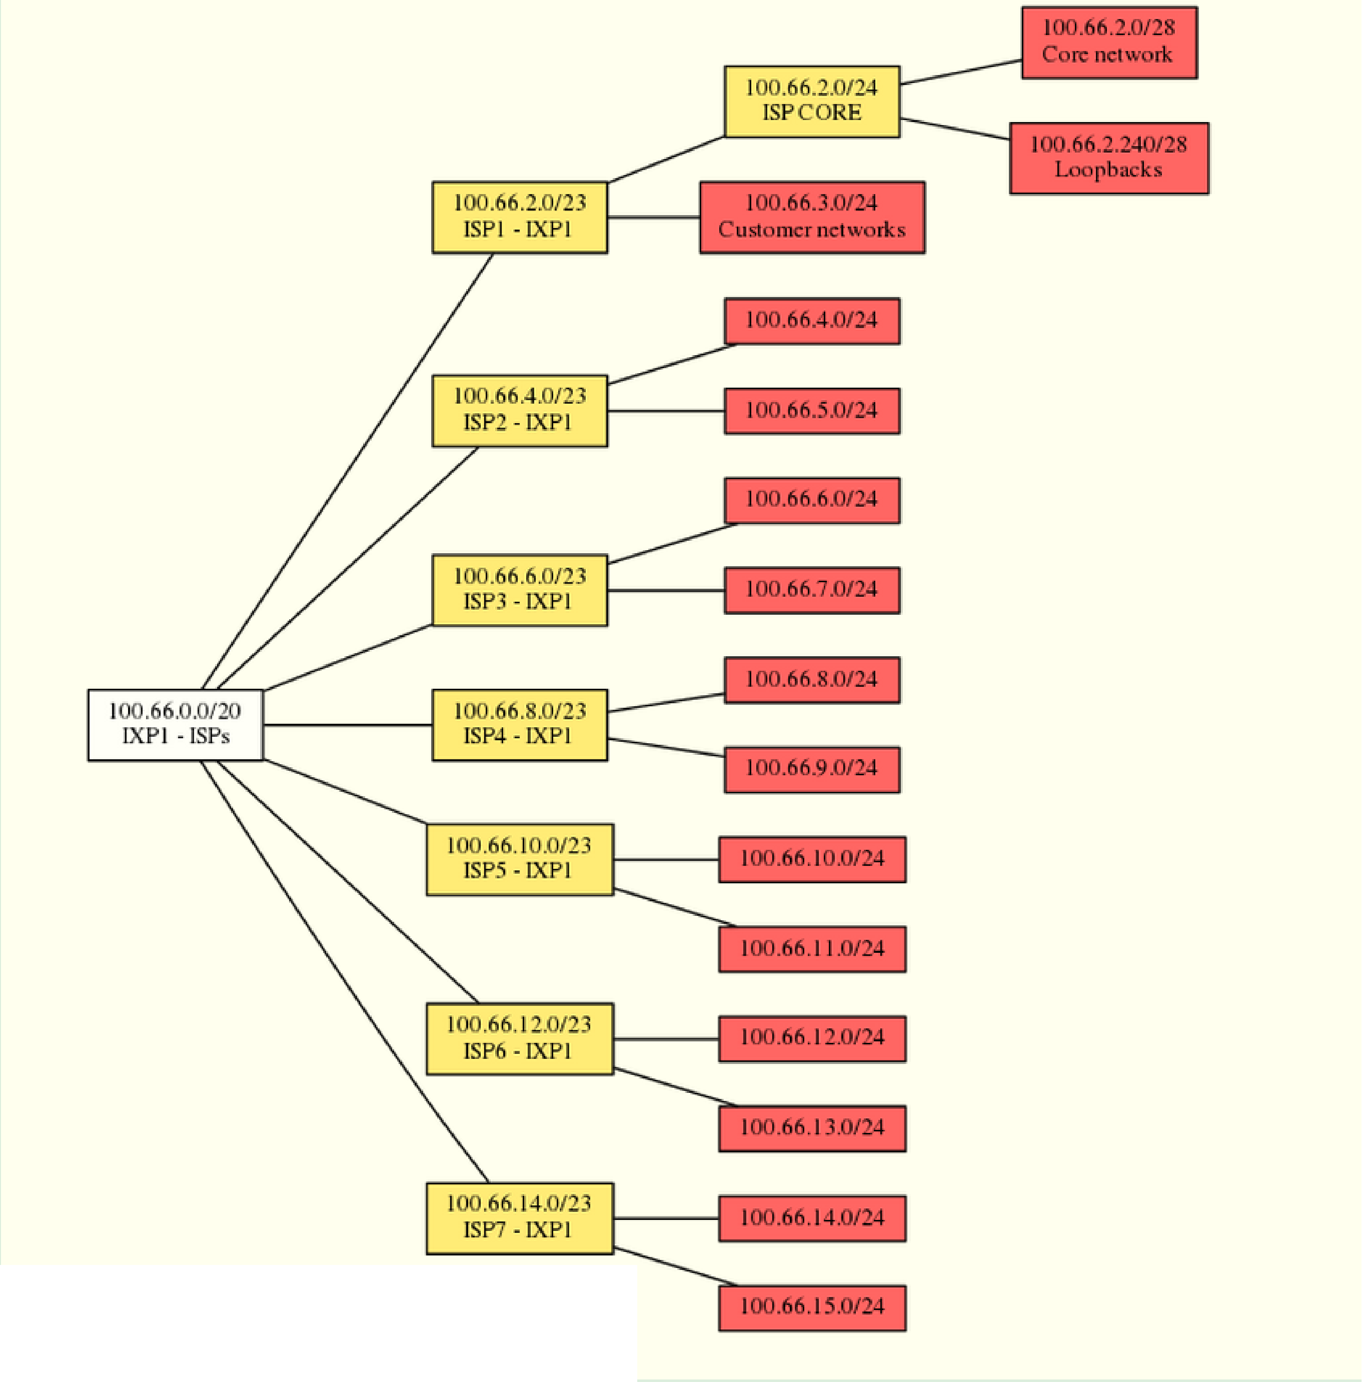

Figure 2 - IXP1 - IPv4 addresses

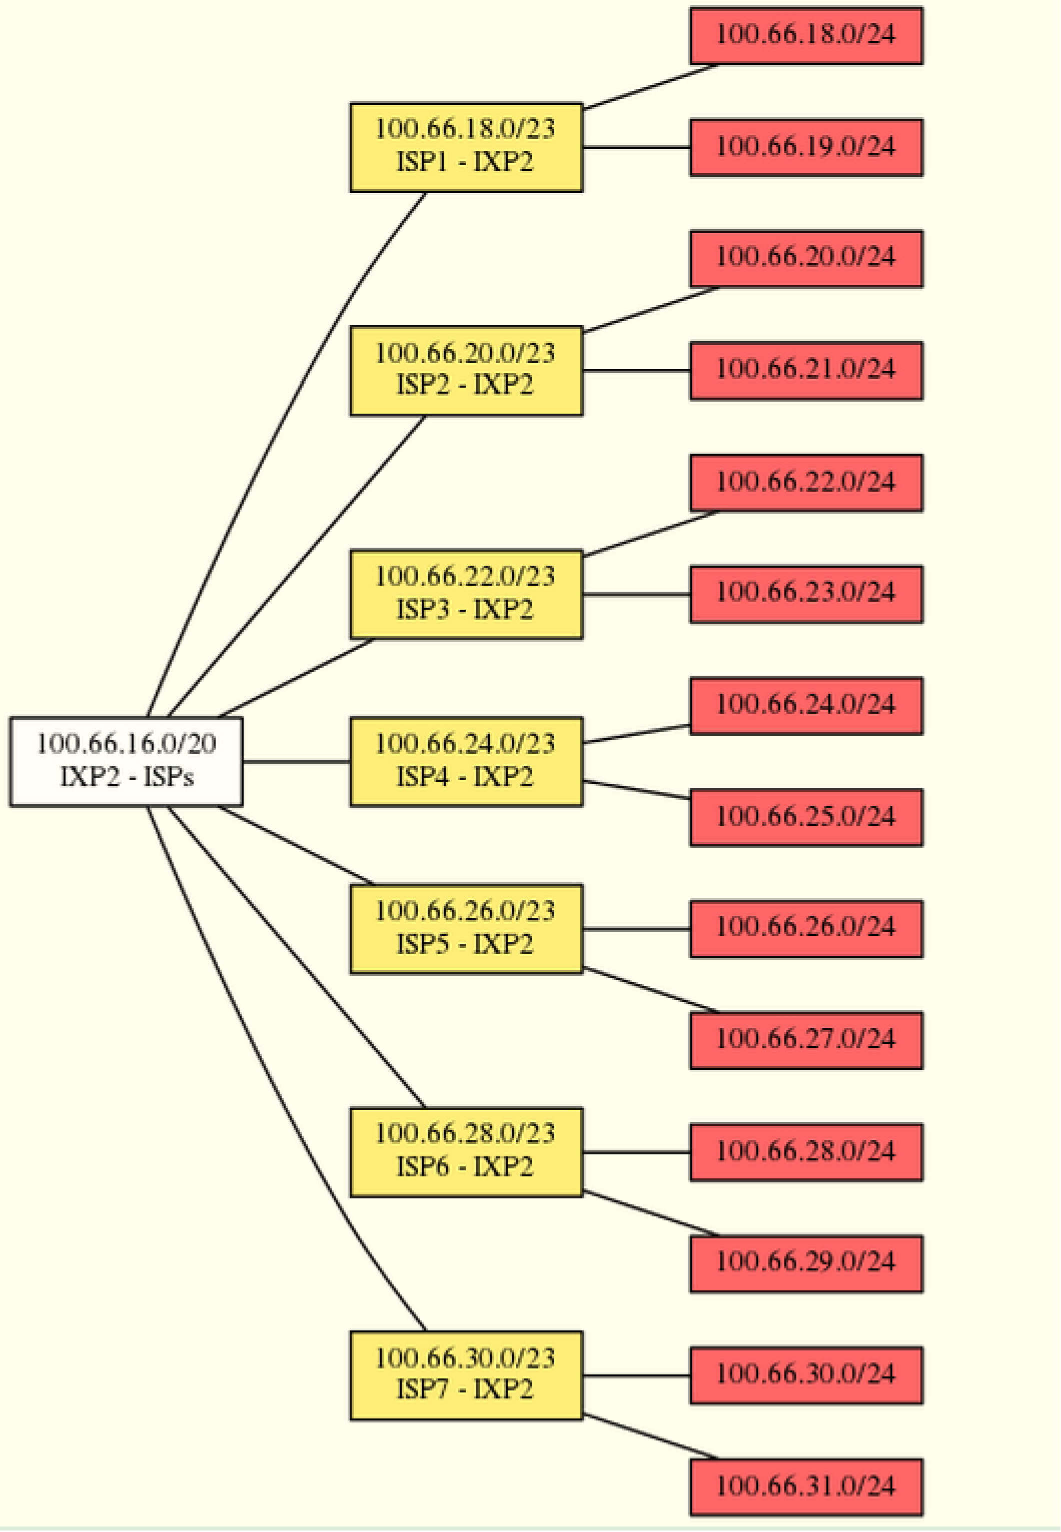

Figure 3 - IXP2 - IPv4 addresses

The typical way that ISPs split up their allocated address space is to carve it into three pieces. One piece is used for assignments to customers, the second piece is used for infrastructure point-to-point links, and the final piece is used for loopback interface addresses for all their backbone routers. The schematic in Figure 2 shows what is typically done.

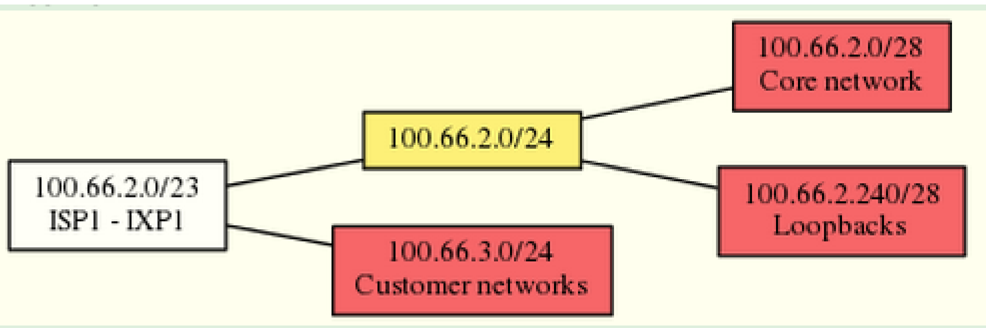

Figure 4 – Dividing allocated block of /23 into Customer, Infrastructure and Loopbacks

ISPs tend to document their addressing plans in flat text files or in spreadsheets – Figure 3 below shows an extract from a typical example (using our addressing scheme here).

Address	Name	Subnet
100.66.2.1	r1-core-isp1-ixp1	100.66.2.0/28
100.66.2.10	r2-core-isp1-ixp1	100.66.2.0/28
100.66.2.251	r1-loop-isp1-ixp1	100.66.2.240/28
100.66.2.252	r2-loop-isp1-ixp1	100.66.2.240/28
100.66.3.1	r2-edge-isp1-ixp1	100.66.3.0/24
100.66.3.10	pc1-edge-isp1-ixp1	100.66.3.0/24

Figure 5 – Extract from an ISP addressing plan

Using the information above, you should now create an address plan for your ISP devices:

Address	Name	Subnet
100.66.		
100.66.		
100.66.		
100.66.		
100.66.		
100.66.		

Figure 6 – ISP addressing plan

6. Ethernet Connections.

The links between the routers are already configured. Note that the Ethernet link in the **Core** of the network has a **/28** mask and the link in the **Edge** or **Customer** network uses a **/24**.

Reminder: *You can check r1-ixp6-isp1 config and r2-isp6-ixp1 config for details. Check that your routers have a similar config.*

7. Ping Test #1.

Ping all physically connected subnets of the neighbouring routers. If the physically connected subnets are unreachable, consult with your neighbouring teams as to what might be wrong. Don’t ignore the problem – it may not go away. Use the following commands to troubleshoot the connection:

```
show arp                               : Shows the Address resolution protocol
show interface <interface> <number> : Interface status and configuration
show ip interface                      : Brief summary of IP interface status
and configuration
```

8. Create Loopback Interfaces.

Loopback interfaces will be used in this workshop for many things. These include generating routes (to be advertised) and configuring some BGP peerings. As discussed earlier in Step 12, we will use part of the allocated IP address block for loopback interfaces. Most ISPs tend to set aside a contiguous block of addresses for use by their router loopbacks.

Our network setup is very simple. There are only two routers in our **ISP** network but in a real ISP network we'd expect to see more. For example, if an ISP had 20 routers, they would need a /27 (or 32 host addresses) to provide a loopback address for each router. We each have 2 routers in our lab – to be prudent and allow for growth, we will set aside a /28 (allows us 16 loopbacks) but only use 2 of them. The assigned loopback addresses are:

Router	Loopback address	Router	Loopback address
r1-isp1-ixp1	100.66.2.251/32	r1-isp1-ixp2	100.66.18.251/32
r2-isp1-ixp1	100.66.2.252/32	r2-isp1-ixp2	100.66.18.252/32
r1-isp2-ixp1	100.66.4.251/32	r1-isp2-ixp2	100.66.20.251/32
r2-isp2-ixp1	100.66.4.252/32	r2-isp2-ixp2	100.66.20.252/32
r1-isp3-ixp1	100.66.6.251/32	r1-isp3-ixp2	100.66.22.251/32
r2-isp3-ixp1	100.66.6.252/32	r2-isp3-ixp2	100.66.22.252/32
r1-isp4-ixp1	100.66.8.251/32	r1-isp4-ixp2	100.66.24.251/32
r2-isp4-ixp1	100.66.8.252/32	r2-isp4-ixp2	100.66.24.252/32
r1-isp5-ixp1	100.66.10.251/32	r1-isp5-ixp2	100.66.26.251/32
r2-isp5-ixp1	100.66.10.252/32	r1-isp5-ixp2	100.66.26.251/32
r1-isp6-ixp1	100.66.12.251/32	r1-isp6-ixp2	100.66.28.251/32
r2-isp6-ixp1	100.66.12.252/32	r2-isp6-ixp2	100.66.28.252/32
r1-isp7-ixp1	100.66.14.251/32	r1-isp7-ixp2	100.66.32.251/32
r2-isp7-ixp1	100.66.14.252/32	r2-isp7-ixp2	100.66.32.252/32

For example, ISP Team 1 on IXP 1 would assign the following address and mask to the loopback on r1-isp1-ixp1:

```
r1-isp1-ixp1(config)#interface loopback 0
r1-isp1-ixp1(config-if)#ip address 104.66.2.251 255.255.255.255
```

Q: Why do we use /32 masks for the loopback interface address?

A: There is no physical network attached to the loopback so there can only be one device there. So we only need to assign a /32 mask – it is a waste of address space to use anything else.

Telnet source address

Most ISPs use the router Loopback address for administrative purposes as well as the anchor point for their network's iBGP sessions. In this step we will configure telnet so that it uses the loopback interface as the source address for all telnet packets originated by the router.

```
ip telnet source-interface loopback 0
```

To check that this has worked, telnet from your router to a neighbouring router and then enter the "who" command. You will see that you are logged in, and the source address will be displayed. For example, using telnet from r2-isp1-ixp1 to r1-isp1-ixp1 gives:

```
r1-isp1-ixp1>who
    Line       User       Host(s)              Idle       Location
*  2 vty 0     isplab     idle                 00:00:00 100.66.2.252
```

Checkpoint #1: *call lab assistant to verify the connectivity. Demonstrate that you can ping and telnet to the adjacent routers.*

9. ISIS with one area and one level (level-2) within the same AS

Each router team should enable ISIS on their router, and use workshop as the ISIS ID in the configuration. In this module, we use level-2 in one area (49.0001) and use wide metrics (IOS default is the historical narrow metric and is not considered good practice). The NET should be 49.0001.x.x.x.x.00, where x.x.x.x represents the router loopback IP address. For example, the loopback for r1-isp1-ixp1 is 104.66.2.251 which will make the NSAP address 49.0001.1040.6600.2251.00

```
Router1(config)# router isis workshop
Router1(config-router)#net 49.0001.1030.6600.2251.00
Router1(config-router)#is-type level-2-only
```

Q: Why do you have is-type level-2-only configured? Write your answer here:

Hint: A nice trick for converting the loopback interface address into the NSAP address is to take the loopback address and put the missing leading zeroes in. For example, r1-isp5-ixp1 loopback address is 104.66.10.251; this is rewritten to 104.066.010.251 putting in the missing zeroes (this is known as Binary Coded Decimal (BCD)). Then rather than having the dot after every third character, move it to be after every fourth character. So 104.066.010.25 becomes 1040.6601.0251 - in summary:

```
IP address in BCD format: 104.066.010.251
System ID:                1040.6601.0251
NET:               49.0001.1040.6601.0251.00
```

10. Setting Wide Metrics.

We also set the metric-style to wide. ISIS supports two types of metric, narrow (historic now and not suitable for modern networks) and wide. IOS still defaults to narrow metrics, so we need to enter explicit configuration to change this to wide. We will set the metric style to be wide for level-2:

```
Router1(config)# router isis workshop
Router1(config-router)#metric-style wide level-2
```

11. Activating ISIS on each interface.

Now that the ISIS process is configured, all connected point to point and shared ethernet interfaces need to be configured with ISIS. Else, you will not be able to see network advertisements via ISIS from routers two or more hops away. Here is an example configuration as would be used on r1-isp6-ixp1:

```
r1-isp6-ixp1(config)# interface fastethernet 0/1
r1-isp6-ixp1(config-if)# ip router isis workshop
```

and on r2-isp6-ixp1:

```
r2-isp6-ixp1(config)# interface fastethernet 0/0
r2-isp6-ixp1(config-if)# ip router isis workshop
```

Note: the ISIS ID on the interfaces must be matched with the router’s ISIS ID.

Q. Can you explain why we chose these interfaces on these routers?

12. ISIS Circuit Type and ISIS Metrics.

Now each team needs to set the circuit type and ISIS metric on each physical interface.

The default circuit type is level-1-2 even though the router has been defined to be a level-2-only router.

The default ISIS metric for all interface types is 10. Unlike OSPF in IOS, ISIS has no automatic scheme to convert the interface bandwidth into a metric value. ISPs deploying ISIS have to come up with their own scheme (as in fact many ISPs using OSPF now also do). In the lab we will use metric 2 for the Ethernet interfaces and metric 20 for the Serial interfaces.

Combining the above, gives the example:

```
Router1(config)# interface fastethernet 0/0
Router1(config-if)# isis metric 2 level-2
Router1(config-if)# isis circuit-type level-2-only
```

13. Announcing the Loopback /32

We do not need to set up ISIS adjacencies on the loopback interface as there are no neighbours there, so we mark it as passive:

```
Router1(config)# router isis workshop
Router1(config-router)# passive-interface Loopback0
```

Note that this will tell ISIS to install the loopback interface address in the ISIS RIB. We do NOT need to add an `ip router isis` statement onto the loopback interface itself. This is different from the required OSPF configuration, and often catches many engineers out, especially those who are learning ISIS after gaining experience with OSPF.

14. ISIS Adjacencies.

Enable logging of ISIS adjacency changes. This is so that a notification is generated every time the state of a CLNS neighbor changes, and is useful for debugging purposes. (Note: From IOS 12.4 onwards, log-adjacency-changes is activated by default when ISIS is first configured.)

```
Router1(config)#router isis workshop
Router1(config-router)#log-adjacency-changes
```

15. Avoiding Traffic Blackhole on Reboot.

When a router restarts after being taken out of service, ISIS will start distribute prefixes as soon as adjacencies are established with its neighbours. In the next part of the workshop lab, we will be introducing iBGP. So if a router restarts, ISIS will start up well before the iBGP mesh is re-established. This will result in the router landing in the transit path for traffic, with out the routing table being completed by BGP. There will not be complete routing information on the router, so any transit traffic (from customer to peer or upstream, or vice-versa) will be either dropped, or resulting in packets bouncing back and forth between adjacent routers. To avoid this problem, we require the router to not announce it is availability until the iBGP mesh is up and running. To do this, we have to provide the following command:

```
Router1(config)#router isis workshop
Router1(config-router)#set-overload-bit on-startup wait-for-bgp
```

This sets ISIS' overload bit such that all routes via this router will be marked as unreachable (very high metric) until iBGP is up and running. Once iBGP is running, the prefixes distributed by ISIS will revert to standard metric values, and the router will pass transit traffic as normal.

16. Ping Test #2.

Ping all loopback interfaces. This will ensure the ISIS IGP is connected End-to-End. If there are problems, use the following commands to help determine the problem:

```
show ip route : see if there is a route for the intended destination
show clns neighbor : see a list of CLNS-IS neighbors that the router sees
show clns interface : see if ISIS is configured and see the IS type
show isis database : see ISIS link state database that the router has
learned
show isis rib : see ISIS routes that the router has learned
show isis topology : see the ISIS topology as learned by the router
```

Checkpoint #2: *call lab assistant to verify the connectivity. Save the configuration as it is on the router – use a separate worksheet, or the workspace at the end of this Module. You will require this configuration several times throughout the workshop.*

17. Traceroute to all routers.

Once you can ping all the loopback addreses, try tracing routes to all the routers using traceroute x.x.x.x command. For example, Router Team 1 would type:

```
r1-isp1-ixp1#traceroute 100.66.2.252
Type escape sequence to abort.
Tracing the route to 100.66.2.252
VRF info: (vrf in name/id, vrf out name/id)
  1 100.66.2.10 36 msec *  28 msec
r1-isp1-ixp1#
```

to trace a route to r2.

Q. Why doesn't the path show the remote loopback address in the path?

A. The path to the loopback is via the physical interface on the remote router.

18. Other Features in ISIS.

Review the documentation or use command line help by typing *?* to see other *show* commands and other ISIS configuration features.

Review Questions

1. What IP Protocol does Ping and Traceroute use?
2. Ping the IP address of your neighbour's router (for example 10.0.15.2). Look at the time it took for the ping to complete. Now Ping the IP address of your router on the same segment (for example 10.0.15.1). Look at the time it took to complete a ping. What are the results? Why is there a difference?
3. What IOS show command(s) will display the router's forwarding table?
4. What IOS show command(s) will display the router's ISIS database?