

Lab 1 – Advanced Router Configuration

Objective: Create a basic physical lab interconnection with two autonomous systems. Each AS should use OSPF, iBGP and eBGP appropriately to construct a working network.

The following will be the common topology used.

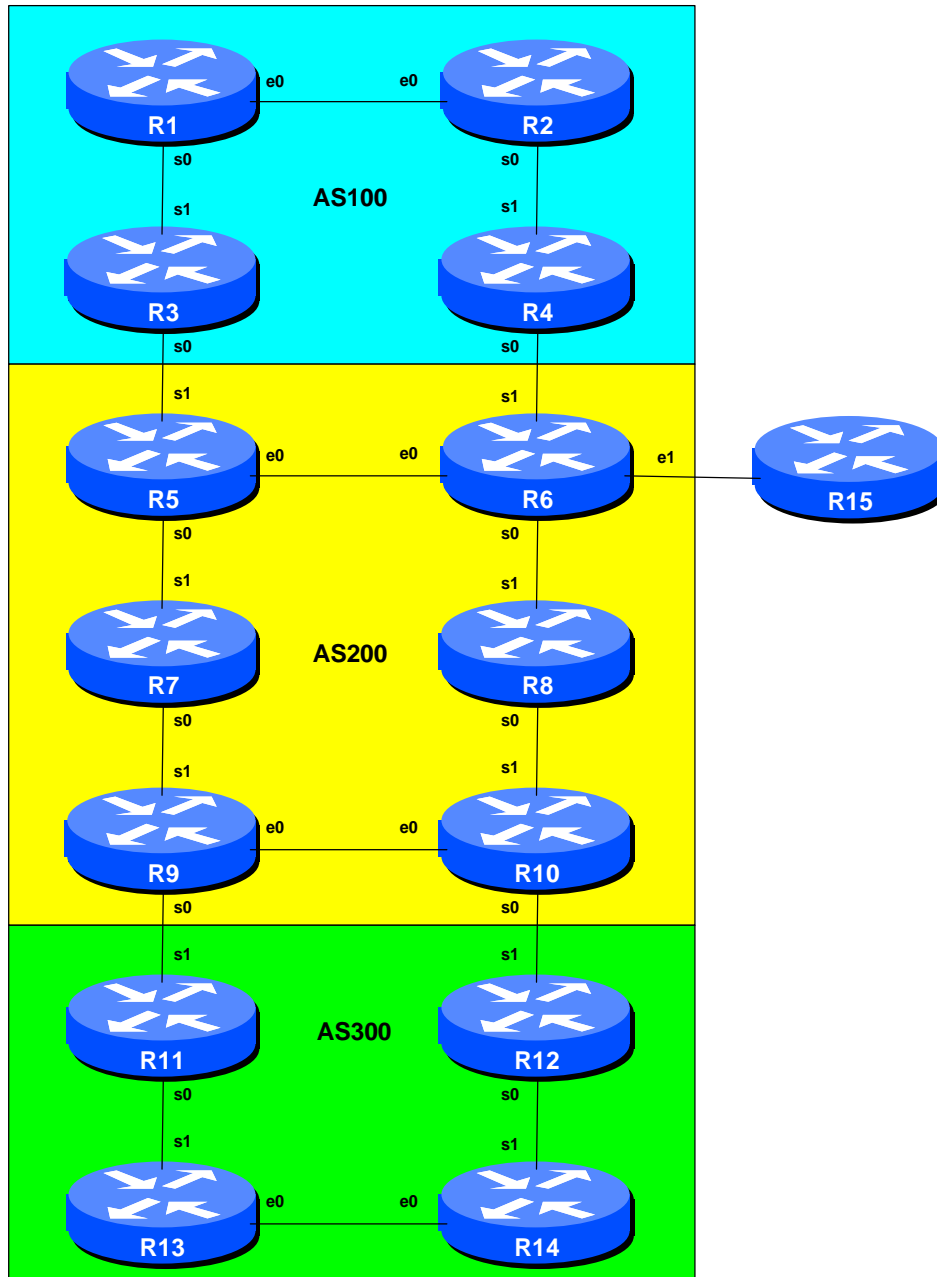


Figure 1 – ISP Lab Basic Configuration

Thursday, July 13, 2000

Lab Notes

The purpose of this module is to construct the workshop lab and introduce participants to IOS 12.0S. It should serve as a reminder on the basic principles of building a network, introducing an IGP, properly function iBGP, and the basics of eBGP:

- 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. Note that Router15 is the Workshop Instructor's router and it will be used at various instances throughout the workshop.
- 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** (OSPF is chosen for this workshop). There is no purpose to building BGP while the chosen IGP (in this case OSPF) is not functioning properly. BGP relies on OSPF to find its neighbours and next hops, and an improperly or non-functioning OSPF 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.
- 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 can usually saves much frustration at a future date or event.

Lab Exercise

The following list is typical for what needs to be done to bring up the lab configuration:

1. **Router Hostname.** Each router will be named according to the table location, Router1, Router2, Router3, etc. Documentation and labs will also refer to *Router1* as R1.

```
hostname Router1
```

2. **Turn Off Domain Name Lookups.** Cisco routers will always try to look up the DNS for any name typed on the command line. You can see this when doing a *trace* on a router with no DNS server or a DNS server with no in-addr.arpa entries for the IP addresses. Unless the Workshop

Instructor specifically tells you that there is a nameserver configured for the lab, we will turn this lookup off for the labs to speed up traceroutes.

```
no ip domain-lookup
```

- 3. Usernames and Passwords.** All router usernames and passwords should be *cisco*. 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.

```
username cisco password cisco
enable secret cisco
service password-encryption
```

- 4. Enabling login access for other teams.** In order to let other teams telnet into your router, you need to configure a password for all virtual terminal lines.

```
aaa new-model
aaa authentication login default local
aaa authentication enable default enable
```

- 5. CIDRise the router.** Make sure the router is configured for CIDR. These two commands are now default in 12.0S, but it is good practice to check just in case:

```
ip subnet-zero
ip classless
```

- 6. Set up timestamps for all logs on the router.** 12.0S has made basic timestamping on the logs the default but ISPs should enable the complete detail on their logs as follows:

```
service timestamps debug datetime localtime show-timezone msec
service timestamps log datetime localtime show-timezone msec
```

Refer to IOS Essentials or the router's on-line help system if you have forgotten what these options mean.

- 7. Enable full logging.** By default on "informational" events are logged and then only to the syslog host. Most ISPs want the full range of log messages to be captured to an internal buffer on the router. This example enables logging into a 64K buffer on the router, capturing the full severity 7 type message (debugging):

```
logging buffered 65535 debugging
```

- 8. Set up a login banner.** Next, set up a login banner. Use an appropriate greeting – consult IOS Essentials document for appropriate and inappropriate greetings. If you use an inappropriate greeting, expect the lab instructors to ask you to change it. For example

Thursday, July 13, 2000

```
login banner ^
Cisco Systems Advanced BGP Workshop Lab
^
```

- 9. Tidy up the vty and console interface configuration.** In the real world we'd now add access-lists to the vty ports on the router. However, this lab is not connected to any external network or the Internet so these will not be required. However, we need to make some changes to the defaults. Basically, only telnet will be the supported mechanism to connect to the routers, and only telnet will be the permitted mechanism to connect from one router to the next. Also, the preferred transport should be changed from *all* to *none*.

```
line con 0
  transport preferred none
line vty 0 4
  transport preferred none
  transport input telnet
  transport output telnet
```

- 10. Create a loopback interface.** Loopback interfaces will be used in this workshop for many things. They are an essential and fundamental requirement for any ISP backbone:

```
interface loopback 0
  description Loopback Interface for RouterXX
```

- 11. Disable pad, finger and bootp servers.** The pad, finger and bootp servers are running by default in IOS. These should be disabled on any Internet router. Finger is a security risk, bootp and pad are simply unnecessary.

```
no service pad
no ip finger
no ip bootp
```

- 12. Remove unneeded SNMP configuration.** IOS installs a default SNMP configuration when the router first starts with an unconfigured NVRAM. As we will not be using SNMP to access the routers in the workshop, this configuration is removed. (Unless configured correctly it is a potential security risk in the Internet.) Example:

```
no snmp-server community public
```

- 13. Time synchronisation.** Router 15 in Figure 1 will always be connected to the workshop network in some way described by the workshop instructors. It's IP address will always be 192.168.1.1/24. And it will be running as a time source with address 192.168.1.2. You should configure ntp to peer with that router, so that the time on your router is synchronised with it and those in the rest of the lab. Don't forget to set the timezone. For example:

```
clock timezone SST 8      ! Singapore Standard Time is GMT+8
```

```
ntp server 192.168.1.2
```

- 14. Enable CEF.** Router platforms from the 2600 upwards support CEF. Those teams with capable hardware should enable CEF now. Note that CEF is not supported or available on the 2500 series.

```
ip cef
```

Entering the *sh cef interface* command will show the status of CEF on the router interfaces. And the command *sh ip cef* will show the forwarding table – currently empty.

- 15. Saving the configuration.** With the basic configuration in place, save the configuration by using “write memory”. Then 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. Don’t forget to frequently save the configuration to NVRAM after each configuration change.

IMPORTANT NOTE: Each router team is strongly recommended to make a copy of the basic router configuration at this stage. It will be assumed throughout this workshop that the above configuration will **ALWAYS** be present on the router. If it is not, each router team will be requested to restore it as a matter of urgency.

Checkpoint #1: call lab assistant to verify the connectivity. Save the configuration as it is on the router either on the worksheet on the end of this hand out, or own your own laptop, or on the classroom tftp server if it is available. It will be required again several times throughout this workshop.

- 16. Back to Back Serial Connections.** Connect the serial connections as in Figure 1. The DCE side of a back to back serial connection is configured with the *clock rate* command that drives the serial circuit.

```
interface serial 0/0
  description DCE Serial Connection to RouterXX
  clock rate 2000000
!
```

- 17. Ethernet Connections.** Use the Catalyst 2924XLs to interconnect the routers using ethernet. *Straight* RJ-45 cables will be used to connect the routers to the switches. The switches will have been configured for multiple VLANs. Ask the instructor if there is doubt as to what the VLAN ranges are.

- 18. IP Addresses.** Each AS is assigned a block of IP addresses.

AS100	220.10.0.0/19	AS300	221.35.0.0/19
AS200	221.19.0.0/19		

Thursday, July 13, 2000


Decide among your team what the addressing plan for you AS should be.

Hint One: point to point links only require /30 blocks.

Hint Two: loopbacks only require a /32 host address.

Hint Three: number your backbone sequentially, from either the start or finish of the range.

Suggestion:



When the IP addresses are assigned, they **MUST** be annotated on the **WHITE-BOARD** at the front of the workshop room. A large network map will have been drawn on the white-board – all the IP address assignments need to be annotated there so that other Router Teams can document and understand the links and routing in this and future modules.

19. 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.

20. OSPF within the same AS. Each router Team should enable OSPF on their router. The OSPF process identifier should be the same as your AS number. All the routers in one AS will be in OSPF Area 0. Use a network statement for every connected interface which will be running OSPF within your autonomous system. Also, use a network statement for the loopback interface address but with a host mask. Don't forget to mark the loopback interface as a passive interface – for example:

```
router ospf 100
 network 221.19.0.0 0.0.0.3 area 0
 network 221.19.0.128 0.0.0.0 area 0
 passive-interface loopback 0
 [ospf] log-adjacency-changes
```

Note: from IOS 12.0 onwards, the command has been changed to drop the “ospf” directive. And from release 12.1 onwards, the “ospf” directive is not supported when enabling “log-adjacency-changes”.

Note: Router6 should also include the network connecting to Router15 in the OSPF announcements, although the interface concerned should be marked as passive. This is so that the network connected to Router15 can be used later in this module.

21. Ping Test #2. Ping all loopback interfaces in your AS. They should all respond. This will ensure the OSPF 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 ip ospf           : see general OSPF information
show ip ospf interface : Check if OSPF is enabled on all intended interface
show ip ospf neighbor  : see a list of OSPF neighbours that the router sees
```

Checkpoint #2: call lab assistant to verify the connectivity. Save the configuration as it is on the router either on the worksheet on the end of this hand out, or own your own laptop, or on the classroom tftp server if it is available.

22. Configuring iBGP Neighbours. Configure iBGP peers within each autonomous system. Use a full iBGP mesh. Don't forget that iBGP peering is configured to be between the loopback interfaces on the routers. Also, it is good practice to use a peer-group. For example:

```
router bgp 100
neighbor ibgp-peers peer-group
neighbor ibgp-peers remote-as 100
neighbor ibgp-peers description iBGP peergroup for internal routers
neighbor ibgp-peers update-source loopback 0
neighbor ibgp-peers send-community
neighbor 221.19.0.129 peer-group ibgp-peers
neighbor 221.19.0.130 peer-group ibgp-peers
neighbor 221.19.0.131 peer-group ibgp-peers
..etc..
```

Use *show ip bgp summary* to check the status of the iBGP neighbour connections. If the iBGP session is not up and/or no updates are being sent, work with the Router Team for that neighbour connection to troubleshoot the problem. Note: get into the habit of using peer-groups and configuring them fully, including the “send-community” directive. This workshop makes extensive use of communities, and making them part of your configuration is good practice.

Q: What are the advantages of using *peer-groups*?

A: BGP peer groups allow a common configuration to be used for several BGP peers. The most common application is for iBGP. All internal BGP peers in an ISP network tend to have the same relationship with each other. Rather than having a substantial configuration per peer, and having to change every peering when details need to be changed, the configuration can be put in a peer-group, and only the peer group has to be changed to alter the peering configuration for all iBGP peers. This substantially reduces the work overhead required in making changes, the router CPU for processing, and significantly cleans up the configuration to view.

Thursday, July 13, 2000

It is advisable that the peer-group command should become the “default” way of configuring all BGP peers. As was stated before, most iBGP peers have the same configuration, so it is of great benefit to the router, the operations staff, and network engineering staff to simply the configurations wherever possible. A configuration which makes extensive use of peer-groups is usually much easier to read than one which has distinct configuration per peer, especially in networks with large numbers of peers.

- 23. Add Prefixes to BGP.** Each Router Team will advertise the CIDR block assigned to them via BGP. AS100 would advertise 220.10.0.0/19, AS200 would advertise 221.19.0.0/19, and AS300 would advertise 221.35.0.0/19:

```
router bgp 100
  no synchronization
  no auto-summary
  bgp log-neighbor-changes
  network 220.10.0.0 mask 255.255.224.0
  !
ip route 220.10.0.0 255.255.224.0 null0 250
```

Don't forget the static route to Null0. This ensures that the prefix has an entry in the routing table, and therefore will appear in the BGP table. Also, don't forget to disable synchronisation and auto-summarisation – these are also mandatory requirements for ISP routers connecting to the Internet. Note that the distance of 250 applied to the static router will ensure that routing protocols announcing this exact prefix will override the static (if this is required/desired).

Checkpoint #3: *call the lab assistant to verify the connectivity.*

- 24. Add an Autonomous System Number pointer.** Even though it is not needed in BGP, the *autonomous-system* command helps people who are troubleshooting the router keep track of the AS number assigned to the router. Example:

```
autonomous-system 100
```

- 25. Enable new format of BGP communities.** It is also worth getting into the habit of changing the BGP community format from the default 32-bit integer to colon separated 16-bit integers, as used in RFC1998. Example:

```
ip bgp-community new-format
```

- 26. Configure eBGP peering.** Now that iBGP is functioning, it is time to configure eBGP. External BGP will be set up between AS100 and AS200, specifically between Routers 3 and 5, and Routers 4 and 6 only. External BGP will also be set up between AS200 and AS300, specifically between Routers 9 and 11, and Routers 10 and 12 only. The remaining lab teams should monitor the BGP table they see on their routers.

Firstly, agree on what IP addresses should be used for the point to point links between the ASes and document them on the whiteboard in the workshop lab. Then configure OSPF so that the DMZ links will be propagated throughout the AS. Don't forget that the DMZ interface should be marked as passive – OSPF will not be running on the DMZ. For example on Router3:

```
router ospf 100
 network 221.19.2.2 0.0.0.3 area 0 ! DMZ link to Router5
 passive-interface serial 0/0
 ospf log-adjacency-changes
```

Once you see the DMZ link appearing in OSPF, configure eBGP between the router pairs, for example:

```
router bgp 100
 neighbor 221.19.2.2 remote-as 200
 neighbor 221.19.2.2 description eBGP with Router5
 neighbor 221.19.2.2 soft-reconfiguration in
```

Use the BGP show commands to ensure that you are receiving prefixes from your neighbouring AS. Don't forget the *soft-reconfiguration* command – this again is mandatory on all eBGP peerings in this workshop. Incremental reconfiguration is “mandatory” on any Internet eBGP router as the disruption through hard resets of the BGP sessions has a severe impact on traffic and customer perceptions of network quality. If you don't remember how soft-reconfiguration works, consult the IOS Essentials documentation and the presentation materials.

Note: There is a new BGP capability called *route-refresh* which is more memory efficient than soft-reconfiguration. If you know the neighbouring router supports this capability, it is a good idea to disable soft-reconfiguration. Cisco routers will attempt to auto-negotiate route-refresh capability at the time the peering is enabled.

- 27. Check the network paths and the routing table.** Run traceroutes between your router and other routers in the classroom. Ensure that all routers are reachable. If any are not, work with the other router teams to establish what might be wrong. Make sure that you can see Router15. The lab instructor will have written the addresses and network up on the whiteboard. (The network is 192.168.1.0/24, the address of the Router6 on that LAN is 192.168.1.254, the address of Router15 is 192.168.1.1.)
- 28. Saving the configuration. For software releases from 12.0 onwards,** the commands to save the configuration are of the format *copy <source> <destination>* where the source and destinations can be any of the following options: *ftp, lex, null, nvram, rcp, running-config, startup-config, system, tftp*. To save the configuration to the TFTP server, use the “*copy system:/running-config tftp:*” command sequence. If the TFTP server is unreachable, “.”s followed by an error message will be displayed rather than “!”s. (Note that the “*write net*” command of earlier releases is still supported but may be removed at a future release.)

Thursday, July 13, 2000

An example of saving the configuration for Router 1 might be:

```
Router1#copy system:/running-config tftp:
Address or name of remote host[]? 192.168.1.4
Destination filename [running-config]? router1-config
!!
2259 bytes copied in 2.920 secs (1129 bytes/sec)
Router1#
```

Checkpoint #4: *call the lab assistant to verify the connectivity.*

29. Summary. This module has covered most of the fundamental configuration topics required to construct an ISP network. It has covered basic router configuration, configuration Best Current Practices, OSPF configuration, iBGP configuration, and finally simple eBGP configuration. No routing policy has been implemented. **Each Router team is strongly recommended to make a copy of their configuration as most of the configuration concepts will be required throughout the remainder of the workshop.**

CONFIGURATION NOTES

Documentation is critical! You should record the configuration at each *Checkpoint*, as well as the configuration at the end of the module.