Module 1 – Basic OSPF and BGP Topology

Objective: Create a basic physical lab interconnection with one OSPF Area and one BGP AS number. Ensure that all routers, interfaces, cables and connections are working properly.

Prerequisites: Knowledge of Cisco router CLI, previous hands on experience.

INSTRUCTOR NOTE:

- Explain room layout and workshop layout to students
- Explain cabling between routers, Catalyst 2924XL, and how the DTE/DCE cables work.
- Emphasise the need to work in teams, and with each other. The workshop is not a competition, but a practical lesson in co-operation to set up a scalable and sustainable network.
- Ask the class to install the Documentation CD in their computer's CD-ROM drives. Make sure that a web browser is available, and that the CD is can be browsed – encourage the class to refer to the CD Documentation whenever they have questions, or are seeking more information, or when directed to by these lab notes. (The Essential Router Configuration Module describes more about an out-of-band network which can be used for Documentation CD access as well.)
- Draw the basic network on the whiteboard before the lab starts – insist that each router team annotate their chosen IP addresses on the board. While many hints are given (see the document describing the Addressing Plan for Modules 1 to 10), it is still important to emphasise the purpose of documenting what being configured.

STEP 5:

INSTRUCTOR NOTE:

Explain the concept of all-zero and all-one subnet in the old classful networking world.

Q: Why is 0/0 a default route?

A: Explain the IP destination address look up process in the routing table, i.e. a comparison between (destination IP address && net mask) and (network prefix && net mask) for all entries in the routing table, then the **longest-matched** entry is chosen.

STEP 11:

INSTRUCTOR TIP:

The concept of a network block reserved for loopback interfaces is usually hard for most students to grasp. Time will need to be spent explaining the concept. The diagram given should help.

STEP 12:

INSTRUCTOR NOTE:

- Explain how the *network* command is used to configure different routing protocols. OSPF uses the command as a way to inform OSPF on which interface it is to run, i.e. start sending hellos

and forming adjacencies. The syntax is *network <network-no.> <inverse-mask> area <area no.>*, where all interfaces in the router whose IP address falls in the range covered by network-no. and inverse-mask will be put in OSPF area <area-no>.

- Then explain why *redistribute connected* command is still needed. It is different from the *network* command in that *redistributed* command will **import** routes into OSPF, i.e. take routes from other sources such as static routes or RIP routes and advertise them into OSPF as external routes.

Checkpoint #1:
INSTRUCTOR TIP1: Make sure that the class save this configuration as it stands. It will be used several times again through out the workshop. If a TFTP server is available, save the configurations there. Explain how to use *"write net"* and why it is good policy to always keep an copy of the current running configuration "off-site".

INSTRUCTOR TIP2: Explain how to use the configuration command *"alias exec <name> [full command]"* to create an alias for frequently used commands. For example, *"alias exec sb show ip bgp"*.

STEP 15:
INSTRUCTOR NOTE:

- Explain why iBGP peering using loopback addresses is a GOOD THING.
  Answer: Since a loopback interface never goes down even when some interfaces go down, the peering would still be up if there exists a path between the peering routers.
- Use every opportunity to explain why loopback interfaces are so useful in IOS. If in doubt, the IOS Essentials presentation should more than convince any class.

STEP 17:
INSTRUCTOR NOTE:

- *Network* command under *router bgp* indicates which network prefix should get advertised in BGP update. However, that network prefix **has to** exist in the IP routing table before it can be advertised by BGP. There is a general rule that the router will never advertise any route it doesn't itself use.
- The network learned from the BGP neighbour is **not** installed in the routing table because BGP is still synchronised with OSPF. The network does not appear in OSPF, so it is not used by BGP, although it is learned from an iBGP neighbour. BGP's *no synchronization* command is introduce in Module 2 – this decouples BGP from OSPF, and if applied in this case will result in the BGP learned networks being installed in the routing table.

STEP 19:
INSTRUCTOR NOTE:

Cisco Systems Inc
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

- Connect the instructor laptop to the console of one of the routers, or to the Cat1900 switch, and demonstrate the routing table and the various "show" commands to the class. Full iBGP peering should be up, so a few "show ip bgp neighbor x.x.x.x route" commands should list which routes are being advertised by which router. It is generally useful to do a traceroute to show the path taken by ip packets to reach a destination.
- Before ending the module, please ask the class to do a "write memory" command so that the configuration is saved to NVRAM. While this is not essential, it saves reapplying some elements of the configuration for the next module.
- If a local TFTP server is available (e.g. instructor laptop), please ask the class to "*write net*" so that each configuration is saved offline. Explain the purpose and good practice to the class. At the end of each module make sure that the class save the router configuration to the local TFTP server.

Module 2 – Basic BGP Configuration Lab

Objective: Simulate four different ISP backbones using a combination of OSPF, internal BGP, and external BGP.

Prerequisites: Module 1

INSTRUCTOR NOTE:

- Before starting the module, explain the aims, and the process being taken to reach a properly working network. It is important to build a working physical network, before installing OSPF, and before attempting iBGP, or eBGP. This maps into real life – one ISP with a bad configuration affects more than just that ISP – it influences the others around him too. Instructors and assistants should ensure that each step in the module is completed properly before attempting the next.

Checkpoint 2:

INSTRUCTOR NOTE:

- Explain the different usage of ***network*** and ***aggregate*** command. The ***network*** command is used to tell BGP which network it should advertise to its peers. ***Aggregate-address*** is used to summarise all the networks that BGP is advertising. This, before ***aggregate-address*** can do its work, the ***network*** command needs to be in place. Don't forget that the network or a subnet must exist in the routing table before the aggregate will be announced!

- Remind the class to save the configuration to NVRAM and (optionally) to the TFTP server before they complete the module – it will be used in the next module.

Module 3 – BGP route filtering and advanced features

Objective: Using the network configured in Module 2, use various configuration methods on BGP peerings to demonstrate neighbour filtering and more advanced IOS features.

Prerequisite: Module 2

INSTRUCTOR NOTE:

- Before starting this module make sure that the configuration from Module 2 is saved. It will be used here. The *aggregate-address* part of the BGP configuration should be removed.
- Explain the purpose of the module before it is started. It is designed to familiarise students with techniques for applying BGP route filtering between peers, not provide a sophisticated fully operational multi-AS "Internet". There will be connectivity breaks between non-connected AS's – explain that to the class. Also, give them hints about what access-lists are required.
- Make it clear to the class that everyone **MUST** complete each step before the class can move on to the next step – otherwise this module does not work.

Checkpoint #1:
INSTRUCTOR NOTE:
The typical answer is the following configuration (for Router1). Note that Router's 2 and 9 have two external peers, in different ASes, so require two different inbound access-lists.

```
router bgp 200
  no synchronization
  network 200.200.4.0 mask 255.255.252.0
  neighbor 200.200.6.2 remote-as 222
  neighbor 200.200.6.2 distribute-list 2 out
  neighbor 200.200.6.3 distribute-list 3 in
  no auto-summary
!
ip route 200.200.4.0 255.255.252.0 null0
!
access-list 2 permit 200.200.0.0 0.0.255.255
access-list 2 deny any
access-list 3 permit 222.222.0.0 0.0.255.255
access-list 3 deny any
```


Cisco Systems Inc
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

Checkpoint #2:
INSTRUCTOR NOTE:

- Explain regular expression symbols and give some examples, such as:

`^200$`	match AS200 only
`.*`	match all ASes starting with the local AS
`^*`	match all ASes
`^$`	match this AS only
`^200_`	match all ASes received from AS200
`_200_`	match all ASes which have AS200 in the path
`_200$`	match all ASes with AS200 origin only, whatever the path
`^200_210$`	match AS210 origin and received from AS200 only
`_200_210_`	match all ASes which have been through AS200 ←→ AS210 link
`^(200_)+$`	match AS200, or AS200 with path stuffing[1]
`^(_[0-9]+)$`	matches one AS, or one AS with path stuffing by the same AS

STEP 7:
INSTRUCTOR NOTE:

- Emphasise that ***neighbor x.x.x.x send-community*** is needed to make BGP send community string along with the prefix in BGP updates.
- Remind the class that those routers one hop away from those with external peerings will not see the community strings set at the borders. This is because the passing of community values has not been set for internal BGP peerings, and communities have not been set for networks originated by local routers.
- Point out to the class why the eBGP peerings only announce the routes which have had the community set. Remember that route-maps have a default deny if the match fails. The match above only sets community on the networks which matches the access-list. As there is no other match close, no other networks are announced in the peering.

Module 4 – Multihoming Strategies Lab

Objective: To demonstrate some of the routing policy and configuration options available when multihoming (or having connectivity to more than one other AS).

Prerequisite: Module 2

STEP 2:
INSTRUCTOR NOTE: Don't make this module more complex than it is. Draw a picture on the board and show the workshop students the **preferred** inbound and outbound paths. Explain why the less specific /16 each AS is announcing becomes the backup path. Supplement Figure 2 with your own drawings and explanation.

Module 5 – Advanced OSPF

Objective: To implement different types of OSPF areas and use advanced OSPF Features in the lab network.

Prerequisite: Module 1 and OSPF presentation

STEP 12:
INSTRUCTOR NOTE: Explain what the 3 NSSA options achieve.

- No-redistribution: do not redistribute any external networks into the NSSA
- Default-information-originate: originate a default route into the NSSA (Type 7)
- No-summary: do not send summary LSAs into the NSSA

STEP 15:
INSTRUCTOR NOTE: There is what may seem to be a confusing situation here. The OSPF virtual link commands "bridge" area 30, yet the OSPF virtual link is actually connecting area 0 with area 40. Router 12 needs the extra configuration line which tells the router that authentication is enabled for area 0. An example of the complete OSPF configuration for Router 12 might as listed below. Check the answer book or CD for complete configurations of all the routers in the workshop.

```
interface serial 0/0
 ip address 222.222.32.2 255.255.255.252
 ip ospf message-digest-key 1 md5 cisco
!
interface serial 0/1
 ip address 220.220.17.2 255.255.255.252
 ip ospf message-digest-key 1 md5 cisco
!
router ospf 100
```

[1] AS path stuffing means seeing an AS path such as 200_200_200 for a network announcement. This is commonly used when defining particular routing policy, for example loadsharing.

Cisco Systems Inc
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
 redistribute connected subnets
 network 220.220.17.0 0.0.0.3 area 30
 network 222.222.32.0 0.0.0.3 area 40
 area 0 authentication message-digest  ! v-link to router7&8
 area 30 authentication message-digest
 area 30 virtual-link 210.210.35.224 message-digest-key 1 md5 cisco
 area 30 virtual-link 220.220.7.224 message-digest-key 1 md5 cisco
 area 40 authentication message-digest
 area 40 nssa no-redistribution default-information-originate no-summ
 area 40 range 222.222.16.0 255.255.252.0
 area 40 range 222.222.32.0 255.255.252.0
 ospf log-adjacency-changes
!
```

Module 6 – Essential Router Configuration

Objective: To configure Out of Band Access, Network Time Protocol and TACACS+ on the ISP Workshop Lab Network.

Prerequisite: Module 1, the IOS Essentials presentation, and Module 5 (optional)

INSTRUCTOR NOTE: Only do steps one to three if a 2511 or router with the async port module is available as not all workshop equipment sets have these routers. If there is no router available which can serve as the OoB system, take time to explain to the class how out of band access works. If there is time, demonstrate OoB access by connecting the auxiliary port one of the classroom routers to the console of another using the supplied console cables.

Module 7 – BGP Route Reflector Lab

Objective: To implement BGP route reflectors as an alternative to fully-meshed iBGP.

Prerequisites: Module 1, the Scaling BGP presentation and Modules 5 & 6 (optional)

INSTRUCTOR NOTE: If continuing from the previous Advanced OSPF Module, Step 1 should be omitted!

Module 8 – Policy Routing

Objective: Using interesting lab exercises, the student will implement some of the fundamental techniques of Policy Routing, as applied in Cisco IOS.

Prerequisite: Module 1 and the Policy Routing Presentation

No Instructor notes.

Module 9 – Internet Exchange Points

Objective: An optional module to demonstrate the use of BGP at Internet Exchange Points.

Prerequisite: Modules 1 to 7 and the BGP presentations

Philip comment: I find this module too hard – Module 19 is more realistic in today's ISP land…

Module 10 – Introduction to Voice over IP

Objective: An optional module to demonstrate the basic configuration of voice over IP on an IP network.

Prerequisites: Module 1

INSTRUCTOR NOTE: A bit of fun at the end of the workshop, if the facilities exist. Probably not worth it unless there are two handsets available in the room.

Module 11 – 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.

Prerequisites: Basic ISP Workshop (at least Modules 1 to 8)

Cisco Systems Inc
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

INSTRUCTOR NOTE:

- These notes assume that this is the first module in the Advanced ISP Workshop. If continuing from the Basic ISP Workshop, less explanation is required.
- Explain room layout and workshop layout to students
- Explain cabling between routers, Catalyst 2924XL, and how the DTE/DCE cables work.
- Emphasise the need to work in teams, and with each other. The workshop is not a competition, but a practical lesson in co-operation to set up a scalable and sustainable network.
- Ask the class to install the Documentation CD in their computer's CD-ROM drives. Make sure that a web browser is available, and that the CD is can be browsed – encourage the class to refer to the CD Documentation whenever they have questions, or are seeking more information, or when directed to by these lab notes. (The Essential Router Configuration Module describes more about an out-of-band network which can be used for Documentation CD access as well.)
- Draw the basic network on the whiteboard before the lab starts – insist that each router team annotate their chosen IP addresses on the board. It is an important process of designing, implementing and operating a network to document what is being configured.

STEP 20:

INSTRUCTOR NOTE:

- The advanced workshop does not use the *redistribute connected subnets* command, but uses the more efficient *network* statement to inject the loopback interface address into OSPF. It may/will be necessary to explain to the class why this is a better method. **Hint**: *redistribute <anything>* is an evil which causes many problems in ISP backbones if not properly managed.

STEP 28:

INSTRUCTOR NOTE:

- Before ending this module (and each module in the workshop), please ask the class to do a "write memory" command so that the configuration is saved to NVRAM. While this is not essential, it saves reapplying some elements of the configuration for the next module.
- If a local TFTP server is available (e.g. instructor laptop), please ask the class to "*write net*" so that each configuration is saved offline. Explain the purpose and good practice to the class. At the end of each module make sure that the class save the router configuration to the local TFTP server.

Module 12 – Multihoming to the Same ISP

Objective: To investigate various methods for multihoming onto the same upstream's backbone

Prerequisites: Module 11 and Multihoming Presentation

No Instructor Notes.

Module 13 – Multihoming to Different ISPs

Objective: To investigate various methods for multihoming onto two different upstream ISPs.

Prerequisites: Module 12 and Multihoming Presentation

No Instructor Notes.

Module 14 – Loadsharing using Communities

Objective: To investigate loadsharing using BGP Communities.

Prerequisites: Module 13 and Multihoming Presentation

No Instructor Notes.

Module 15 – Multihoming to the Same ISP using RFC1998

Objective: To investigate various methods for multihoming onto the same upstream's backbone

Prerequisites: Module 12 and Advanced Communities Presentation

No Instructor Notes.

Module 16 – Multihoming to Different ISPs using RFC1998

Objective: To investigate various methods for multihoming onto two different upstream ISPs.

Prerequisites: Modules 13 and 15, and Advanced Communities Presentation

No Instructor Notes.

Module 17 – Advanced Community Usage (RFC1998++)

Objective: To investigate various methods for multihoming onto two different upstream ISPs .

Prerequisites: Modules 15 and 16, and Advanced Communities Presentation

No Instructor Notes.

Module 18 – Transit

Objective: To investigate methods for providing transit services.

Prerequisites: Modules 12 and 13, and the Transit Presentation

No Instructor Notes.

Module 19 – Internet Exchange Points

Objective: To investigate methods for connecting to an Internet Exchange Point.

Prerequisites: Modules 12, 13 and 18, and the Exchange Points Presentation

No Instructor Notes.

Module 20 – Overseas Collocation

Objective: To investigate methods for connecting to Internet backbones overseas.

Prerequisites: Modules 12, 13, 18 and (optionally) 19, and the Collocation Presentation

No Instructor Notes.

Cisco Systems Inc
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100