

Static Routing Lab

Campus Network Design & Operations Workshop


Introduction

The purpose of this lab is to now add Layer 3 functionality to our newly finished Layer 2 network. The Core Switch has basic L3 capability right now in that it is routing between the different VLANs configured on it. However, we now want to give the campus connectivity to the border router and onwards to the NREN and the Internet.

The class will continue in its 6 groups with each group being assigned one campus to configure; attention will now be turned to configuring routing on the border router and core router (which we have been calling a switch up to now).

Lab Layout

The following diagram shows the layout of the devices and all the links for each campus:



Our campus network consists of two routers, bdr1.campusX and core1.campusX as well as six switches. We have already configured the distribution and edge switches, as well as completed the L2 configuration of the core router. We will not need to touch any of the L2 configuration any more.

Accessing the Lab

The Workshop Instructors will have already told you what the lab environment is. It will either be run on a Virtual Platform, or on real physical switches provided in the Training Room.

Refer to the **correct** document below for information about logging into the devices that have been assigned to your group:

VIRTUAL ENVIRONMENT: [Virtual Environment Lab Access Instructions](#)

PHYSICAL HARDWARE: [Physical Hardware Lab Access Instructions](#)

Border Router Initial Configuration

Hostname

Your border router should be given a basic configuration as follows:

```
Router> enable
Router# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# hostname bdr1.campusX
bdr1.campusX(config)#
```

Turn Off Domain Name Lookups

Cisco devices 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 in-addr.arpa entries for the IP addresses. We will turn this lookup off for the labs for the time being to speed up traceroutes.

```
bdr1.campusX(config)# no ip domain lookup
```

Set the Domain Name

We will now set the domain name of our campus devices, for future use in this workshop.

```
bdr1.campusX(config)# ip domain name ws.nsrc.org
```

Configure console and other ports

```
bdr1.campusX(config)# line con 0
bdr1.campusX(config-line)# transport preferred none
bdr1.campusX(config-line)# exec-timeout 30 0
bdr1.campusX(config-line)# line aux 0
bdr1.campusX(config-line)# transport preferred none
bdr1.campusX(config-line)# exec-timeout 30 0
bdr1.campusX(config-line)# line vty 0 4
bdr1.campusX(config-line)# transport preferred none
bdr1.campusX(config-line)# transport input all
bdr1.campusX(config-line)# exec-timeout 30 0
```

("transport input all" is to allow inbound telnet and ssh. The version of IOSv used on the border router has them disabled by default)

Username and Passwords

All router usernames should be **cnclab** with password being **lab-PW**. The enable password (which takes the operator into configuration mode) needs to be **lab-EN**¹.

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.

```
bdr1.campusX(config)# username cnclab secret lab-PW
bdr1.campusX(config)# enable secret lab-EN
bdr1.campusX(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 secret² which is already encrypted).

Enabling login access for other systems

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.

```
bdr1.campusX(config)# aaa new-model
bdr1.campusX(config)# aaa authentication login default local
bdr1.campusX(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 the operator to check the history of what events happened on the router.

```
bdr1.campusX(config)# no logging console
bdr1.campusX(config)# logging buffered 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 show log should be used at the command prompt.

And we also want to set up improved time-stamping for the log messages as well:

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

which will give resolution down to milliseconds, and include the year as well.

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.

```
bdr1.campusX(config)# end
bdr1.campusX# write memory
Building configuration...
[OK]
bdr1.campusX#
```

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.

Configure interfaces on the Border Router

Configure the interface towards the NREN

The full address plan for the lab can be found in the [IP Address Plan](#). Consult the address plan for the addresses of the point to point links between the Campus Border Router and the NREN Router.

Make sure you change the **X** and **N** below to the correct value from address plan mentioned above:

```
interface GigabitEthernet0/0
description Link to NREN
ip address 100.68.0.N 255.255.255.252
ipv6 address 2001:DB8:100:X::1/127
no ip redirects
no ip proxy-arp
ipv6 nd prefix default no-advertise
ipv6 nd ra suppress all
no shutdown
```

Test that you can ping the NREN end of the link.

Configure the interface towards the Core Router

Make sure you change the **X** below to the correct value for your campus:

```
interface GigabitEthernet0/1
description CAMPUS CORE
ip address 100.68.X.1 255.255.255.240
ipv6 address 2001:DB8:X:0::1/64
no ip redirects
no ip proxy-arp
ipv6 nd prefix default no-advertise
ipv6 nd ra suppress all
no shutdown
```

Configure remaining interfaces on the Core Router

Configure interface towards the Border router

We will now configure the interface on the Core Router connecting it to the Border Router. Since we will never carry multiple VLANs on this link, we can configure the port as a pure layer 3 (routed) interface, using the no switchport command.

Make sure you change the X below to the correct value for your campus:

```
interface GigabitEthernet0/0
description CAMPUS CORE to BORDER
no switchport
ip address 100.68.X.2 255.255.255.240
ipv6 address 2001:DB8:X:0::2/64
no ip redirects
no ip proxy-arp
ipv6 nd prefix default no-advertise
ipv6 nd ra suppress all
no shutdown
```

Test that you can ping your Border router at the other end this link.

Configure the Network Management and Monitoring interface

Our network management and monitoring server, srv1.campusX.ws.nsrc.org, is connected to GigabitEthernet0/3 on the core router. We'll configure the router, core1.campusX, so that we can start to use that server to manage and monitor our network. Again, since we are only routing a single subnet on this interface, we can configure the port as a routed interface using the no switchport command:

```
interface GigabitEthernet0/3
description Network Management and Monitoring
no switchport
ip address 100.68.X.129 255.255.255.240
ipv6 address 2001:DB8:X:1::1/64
no ip redirects
no ip proxy-arp
ipv6 nd prefix default no-advertise
```

```
ipv6 nd ra suppress all
no shutdown
```

Now connect to the console port of srv1 in your campus. Login as the "sysadm" user, with the password given to you by the workshop instructors.

From srv1, you should be able to ping the core router using the addresses you've just added to its GigabitEthernet0/3 interface:

```
ping 100.68.X.129
ping6 2001:DB8:X:1::1
```

If that works, try using telnet to connect to the router.

```
telnet 100.68.X.129
```

To disconnect, press ctrl-] then type close.

We are doing this just to verify connectivity; we will come back and make use of the network management and monitoring server later on in the workshop.

Configure IP addresses for the STAFF and STUDENT VLANs.

We now configure the IP addresses of the STAFF and STUDENT VLANs on the **Core Router**. Here is an example for VLAN 11, the STAFF VLAN in Building 1:

```
interface vlan 11
description STAFF VLAN of Building 1
ip address 172.2X.11.1 255.255.255.0
ipv6 address 2001:DB8:X:11::1/64
no shutdown
```

Replace **X** with your campus number. Do the same for all the other VLANs as well.

Once completed, your Core switch should have VLANs 11, 12, 21 and 22 configured with both IPv4 and IPv6 addresses.

Question: What is significant about the IP address configured on each VLAN?

Answer: This IP address now serves as the default gateway for every device connecting to this VLAN. So a user on VLAN 11 will have default gateway 172.2X.11.1, etc.

NB: This configuration only goes on the core switch, not the distribution and access switches.

Configure Static Routing

At this stage you should be able to ping each of the devices in your campus network from their immediate neighbours. If you try to ping the Border router from one of the switches or the NMM server you'll have less success. We need to add some additional routing information to the routers so that we can pass packets successfully.

Let's look at the routing information on the Core router:

core1.campus1#show ip route

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2

ia - IS-IS inter area, * - candidate default, U - per-user static route

o - ODR, P - periodic downloaded static route, H - NHRP, I - LISP

a - application route

+ - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is not set

100.0.0.0/8 is variably subnetted, 4 subnets, 2 masks

C 100.68.1.0/28 is directly connected, GigabitEthernet0/0

L 100.68.1.2/32 is directly connected, GigabitEthernet0/0

C 100.68.1.128/28 is directly connected, GigabitEthernet0/3

L 100.68.1.129/32 is directly connected, GigabitEthernet0/3

172.21.0.0/16 is variably subnetted, 12 subnets, 2 masks

C 172.21.10.0/24 is directly connected, Vlan10

L 172.21.10.1/32 is directly connected, Vlan10

C 172.21.11.0/24 is directly connected, Vlan11

L 172.21.11.1/32 is directly connected, Vlan11

C 172.21.12.0/24 is directly connected, Vlan12

L 172.21.12.1/32 is directly connected, Vlan12

C 172.21.20.0/24 is directly connected, Vlan20

L 172.21.20.1/32 is directly connected, Vlan20

C 172.21.21.0/24 is directly connected, Vlan21

L 172.21.21.1/32 is directly connected, Vlan21

C 172.21.22.0/24 is directly connected, Vlan22

L 172.21.22.1/32 is directly connected, Vlan22

and on the Border router:

bdr1.campus1#show ip route

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2

ia - IS-IS inter area, * - candidate default, U - per-user static route

o - ODR, P - periodic downloaded static route, H - NHRP, I - LISP

a - application route

+ - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is not set

100.0.0.0/8 is variably subnetted, 4 subnets, 3 masks

```
C    100.68.0.0/30 is directly connected, GigabitEthernet0/0
L    100.68.0.2/32 is directly connected, GigabitEthernet0/0
C    100.68.1.0/28 is directly connected, GigabitEthernet0/1
L    100.68.1.1/32 is directly connected, GigabitEthernet0/1
```

Each of the routers knows about the **local** and **connected** networks but no other routes.

What about IPv6? (Use show ipv6 route). What routes do you see for IPv6 destinations? Is there a similarity with what you see for IPv4?

Turn on IPv6 Routing on Core and Border Routers

Cisco IOS routers have IPv6 Routing turned off by default. So while we can reach our directly attached neighbours, we cannot get anywhere beyond, nor can we turn on any IPv6 routing protocols. We now need to turn on IPv6 routing, and to do that we use the following two commands:

```
ipv6 unicast-routing
ipv6 cef
```

(IPv4 routing is of course already on by default. At some point in the future Cisco may well turn on IPv6 routing by default.)

Static routes on the Core router

The Core router needs a default route added to it so that we can forward traffic from the Campus network to the wider Internet via the NREN. We add this route to send traffic to the border router:

```
ip route 0.0.0.0 0.0.0.0 100.68.X.1
ipv6 route ::/0 2001:DB8:X:0::1
```

Static routes on the Border router

The Border router needs a default route added to it so that we can forward traffic from the Campus network to the wider Internet via the NREN. We add this route to send traffic to the NREN router:

```
ip route 0.0.0.0 0.0.0.0 100.68.0.N
ipv6 route ::/0 2001:DB8:100:X::0
```

Choose the correct value for **X** and **N** from the IP address tables we used when we set up the interface.

You have also added a number of subnets on your core router and building switches for the Network Monitoring and Management subnet and for VLANs 10, 11, 12, 20, 21 and 22. Your Border router needs to be able to send packets to those subnets.

Here we are adding the static routes for Building 1's VLANs pointing to the core router:


```
ip route 172.2X.10.0 255.255.255.0 100.68.X.2
ip route 172.2X.11.0 255.255.255.0 100.68.X.2
ip route 172.2X.12.0 255.255.255.0 100.68.X.2
```

Do the same for Building 2 routes:

```
ip route 172.2X.20.0 255.255.255.0 100.68.X.2
ip route 172.2X.21.0 255.255.255.0 100.68.X.2
ip route 172.2X.22.0 255.255.255.0 100.68.X.2
```

And now add the IPv6 static routes on the border router for Building 1:

```
ipv6 route 2001:DB8:X:10::/64 2001:DB8:X:0::2
ipv6 route 2001:DB8:X:11::/64 2001:DB8:X:0::2
ipv6 route 2001:DB8:X:12::/64 2001:DB8:X:0::2
```

and for Building 2:

```
ipv6 route 2001:DB8:X:20::/64 2001:DB8:X:0::2
ipv6 route 2001:DB8:X:21::/64 2001:DB8:X:0::2
ipv6 route 2001:DB8:X:22::/64 2001:DB8:X:0::2
```

Also add in the IPv4 and IPv6 static routes for the Server network:

```
ip route 100.68.X.128 255.255.255.240 100.68.X.2
ipv6 route 2001:DB8:X:1::/64 2001:DB8:X:0::2
```

Set up the default gateway on the distribution and edge switches

The switches also need a default route added to them so that their MGMT VLAN can forward traffic to elsewhere in the network - without it, management traffic will only ever be able to reach other devices on their own VLAN.

On each switch we add this route to forward traffic to the Core router.

For Building 1 try:

```
ip route 0.0.0.0 0.0.0.0 172.2X.10.1
ipv6 route ::/0 2001:DB8:X:10::1
```

For Building 2 try:

```
ip route 0.0.0.0 0.0.0.0 172.2X.20.1
ipv6 route ::/0 2001:DB8:X:20::1
```

Testing the routing setup

Now try pinging 8.8.8.8 from your Core router - does this work? And from srv1?

What about a traceroute to 8.8.8.8? What do you see?

You should be able to see the trace go to your Border router, to the NREN router you connect to, to the lab gateway, and then on out to the Internet. The actual addresses (and names) will of course depend on the network operator's infrastructure the lab has been connected to.

Now try pinging 2001:4860:4860::8888 from your Core router. Try a traceroute as well.

What happens? Can you explain why the traceroute stops displaying anything³ after 2-3 hops?

Checkpoint: call an instructor and show them your working system.

1. 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 network operator ever use easily guessable passwords as these on their live operational network. 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.[↵](#)
2. for IOS releases prior to 12.3, the username/secret pair was not available, and operators would have had to configure username/password instead. Do **NOT** use the username/password combination, nor the "*enable password*" directive - these use type-7 encryption which is not secure at all, whereas the "*secret*" uses the more secure md5 based encryption.[↵](#)
3. To interrupt a traceroute (or other command line instruction) that has "hung", use the Cisco break sequence CTRL-SHIFT-6.[↵](#)