JANET
IPv6 Hands-on Workshop

Lab 1: **First Look**

UKERNA, Lancaster University
and University of Southampton, 2006

Laboratory Overview

- Enable IPv6 on a node
- Enable IPv6 on subnets
- Static, default routing
- Network administration tasks (diagnostics & tools)

These exercises offer an overview of IPv6 in operation from both a network operator's perspective and that of a workgroup administrator.

All the commands throughout the lab sessions have been provided for both Cisco's IOS and Juniper's JUNOS router environments. For this workshop you'll only need to follow the commands in the sections labelled JUNIPER and BOTH. We included the IOS equivalents since most attendees stated they have familiarity with IOS.

The steps that you will take in this laboratory will be to enable and configure (in as much as there is to configure) IPv6 on workstations; configure a workgroup router to participate in an IPv6-over-Ethernet Enterprise network with static routing.

Places to look when seeking on-line documentation include:

* Cisco Software IOS Releases Command References
  http://www.cisco.com/univercd/cc/td/doc/product/software/

•Juniper Junos Software Documentation
  http://www.juniper.net/techpubs/software/junos/

Both of these documents (as well as other resources for this workshop) are linked from
  http://www.ipv6.org.uk/workshop/

Note that you're not expected to finish everything in this lab - it's a first look for you to experiment at your own pace.

IOS / JUNOS

- Knowledge of Cisco IOS / JUNOS is useful.
- If new…
  - All administration takes place under configuration mode.

Operation	IOS	JUNOS
Connection Type	Telnet (Single User)	ssh (Multi User)
Show Command	show ?	show ?
Enter Administration Mode	enable (requires password)	**user dependant*
Enter Configuration Mode	config terminal	configure

The workgroup routers start off in a state in which you can access their Command Line Interface (CLI) using ssh/telnet. The username (ssh only) is "ukwuser" and the password is "ukw$password".

CISCO

In Cisco IOS there is not much which can be performed at this level and level-15 access is required, this can be accessed by the "enable" command and entering the password "ukw$enable".

JUNIPER

In JUNOS the different levels of access are granted per user, thus operations at this level can vary. The ukwuser account was created from the terminal only root account and has admin (the highest level) access. Note that you can change the terminal type to enable some commands to execute properly e.g. "monitor interface traffic" requires a terminal type to display properly, this can be set with the "set cli terminal ?" command and choosing one of the options in place of the "?". In your case this is not necessary.

BOTH

At this level we can execute a few of the basic commands to get used to the interface:-

"show ?" lists a set of possible command completions for the current context. Depending on your current privilege level, different options will avail themselves to you. Try IOS ["show ip interfaces"] JUNOS["show interfaces"] and press return, and have a look over the output.

CONFIGURATION MODE

To enter configuration mode enter the relevant command shown on the slide above. Note how the prompt changes to reflect the new context - currently the router's global configuration context. Additional sub-contexts include interface definition mode where subsequent directives entered affect a particular interface; line definition mode to configure the administration and modem interfaces; etc. To step back up a level to the parent context (and from the global context back out of configuration mode and back into exec/interactive), type "exit" on a line of its own.

Configurations

- You should know how to save configuration changes
- If you get ‘stuck’ then please speak to a helper - your configuration can then be reset to its state as per the start of the lab session
- Before each session, we will preconfigure the routers accordingly

Writing Configuration Changes

CISCO

Once you’ve made configuration changes, the live running configuration can be dumped to console by typing “show running-config” in interactive mode. Likewise the configuration directives that will be applied post reboot is the ‘startup-config’. To make the current live config persist (I.e. become the startup-config), enter “write memory” in interactive mode.

JUNIPER

Juniper makes a candidate configuration which you edit and when satisfied you can “commit” this configuration. Commit also has many options accessible in the usual way “commit ?”. To commit changes you are happy with type “commit” which will check the configuration and then commit it. “commit confirmed” applies the configuration for 10 minutes and will then roll back unless confirmed with a “commit”. It is also possible specify the amount of time before a commit rolls back.

CRASH COURSES

Courses in the two very similar systems is beyond the scope of our workshop however there are some excellent online resources available which are not too strenuous.

Juniper Free Video Course

http://www.juniper.net/training/elearning/junos_cli/index.html

Your cluster

- Each group has the following:
  - Juniper J4300 Router
    - Four Ethernet Ports
  - Client nodes
    - One Windows XP Service Pack 2 PC
    - One Fedora Core 5 Linux PC
    - … split into two subnets
  - Full IPv4 connectivity

Our hosts, ISS at Lancaster, have kindly delegated an entire Class C IPv4 network for this workshop, resulting in every workstation and router having public and globally-routable IPv4 addresses.

The laboratory is split into six groups (A through F), with each group sharing a bench comprising two workstations, each on their own subnet. One workstation runs Windows XP, the other Linux.

Each bench has an A3 network topology map showing the initial state of the laboratory network.

You will have full administrator privileges on the workstations and on your local workgroup router.

The windows administrator account is called “admin”, and the Linux one is called “root”. The password for both is “ukw$password” .

Note that there is only one router per group, so you will need to work together and rotate responsibility for configuring the router, etc.

The J4300 can be configured by more than one person at the same time however this situation can be avoided with the “configure exclusive” command. When leaving the configuration mode, if changes have been made without committing, you will get a warning telling you such.

Current setup

- IPv6 in ‘provider’
  - Routing IPv6 towards your edge router (from the access router)
  - **No** edge router interfaces configured for IPv6.
  - **No** edge router IPv6 configuration toward client subnets.
- Stock IPv4 on your client subnets
  - Need to install IPv6 and witness Stateless Address Autoconfiguration (SLAAC) at work.
- Next step: get the clients ready

It is probably worth spending a little time getting familiar with the topology of the laboratory network.

The groups have been allocated so that you should all have at least one experienced operator/administrator per group, so if you’re not familiar with core tools such as tcpdump, Ethereal, ping, traceroute, then please do ask.

Ex1. IPv6 on client nodes

- IPv6-enable the clients first
  - In operational environments you would most likely get routing infrastructure ready first and then start advertising prefixes for SLAAC, or enabling a stateful-DHCP service
- Client Configuration
  - Observe network configuration, interface statuses, etc.
  - Install/enable IPv6
  - Observe changes
  - Watch bits on the wire as a result of enabling IPv6
    - SLAAC, DAD, ND, etc.
  - Observe ping tests on a local link

For the purposes of this laboratory, you will IPv6 enable the two clients to see what addresses are auto-configured without any routing infrastructure available, and to then observe the effects of introducing IPv6 Router Advertisements and global addresses on a link.

In a production environment, we recommend that you first configure the routing infrastructure, with the global prefixes ubiquitously routable within your enterprise, but *without* advertising prefixes in router advertisements, i.e. have the infrastructure in place (and secured) before client nodes start trying to use it.

We recommend you start Ethereal/tcpdump now, listening for IPv6 on your primary interface, we shall be referring to the output of these tools throughout the lab.

TCPDUMP

tcpdump -i eth0 ip6

ETHEREAL

Capture Menu -> Options

Set the Interface to capture from

Set the capture filter to “ip6”

Click Start

Ex1a. IPv6 in XPsp2

- Observe current, non-IPv6 setup
- Run Ethereal
- Enable IPv6 on the interface
- Observe interfaces, routes
- ping6 all-nodes multicast
- Observe Ethereal results

All the Windows XP machines in the lab should be currently configured to be IPv4-only. This exercise will enable IPv6 and show its initialisation.

ipconfig (/all) provides information on interface IP addresses under Windows

"route print" shows the routing table.

Enabling IPv6 on the Interface can be done via Control Panel / Network Connections / Local Area Connection / Properties. A new Protocol of "Microsoft TCP/IP version 6" can then be installed. Alternatively this can be done from the command prompt by typing ipv6 install. Various tunnel adapter interfaces may appear in ipconfig at this point, however these are not in use as there will be native IPv6 on the link.

Better information about what is configured can be found using "netsh". Run "netsh" on a command prompt and you will be given a further netsh prompt. Information can be found under the "interface" context. Type "interface ip" to go to the IPv4 information. "help" lists commands, "show address" may be of interest. ".." moves up a context. "interface ipv6" contains very detailed information about the IPv6 configuration. Browse around the available commands, "show address", "show routes" and "show neighbors" are good starting points.

Running an ipconfig should now reveal that you have a link-local fe80:: IPv6 address, however no global prefixes have been discovered that this point.

Ex1b. IPv6 in FC 5

- Observe current setup / install IPv6
  - ifconfig
  - 2.6+ kernels come with IPv6 enabled.
- Check interfaces, routes

You can use the standard ‘ifconfig’ and ‘route’ tools to view system configuration. Alternatively, ‘ip addr show’ and ‘ip route show’ (for IPv4) and ‘ip -6 route show’ (for IPv6) can be used, if you are more familiar with the ‘ip’ tool (and it is installed on the Linux machine in question as it is at the lab).

IPv6 support is provided by the ‘ipv6’ kernel module, which comes built into the 2.6+ kernel we are using here. Older 2.4 kernels require this to be enabled usually by adding the line NETWORKING_IPV6=yes to /etc/sysconfig/network, or the kernel configuration itself (requires recompiling). IPv6 can also be built as a module in older 2.4 kernels and inserted on startup.

With IPv6 enabled on both machines, using link-local addressing you could multicast ping (ping6 -I eth0 ff02::1) other clients on your subnet, however since there are no other clients and the router is not yet configured no responses would be seen apart from your own.

Ex2. IPv6 on workgroup router

Four steps:

1. Global router configuration
2. Interface facing the upstream (fe-0/0/0)
   - Enable protocol
   - Configure address
3. Interfaces facing subnets (fe-1/0/X)
4. Static routing

Ex2a. Global config

- CISCO
  - Enable IPv6 datagram forwarding
  - Enable IPv6 CEF (optional)
  - Configure an IPv6-transport name server (optional)
- JUNIPER
  - IPv6 is available for use and needs no global configuration.

CISCO

There are two important directives that need to be configured on the Ciscos to enable them to route IPv6. First, “ipv6 unicast-routing” will enable the forwarding of IPv6 unicast datagrams. Secondly, if you are running Cisco’s Expedited Forwarding protocol for IPv4, we need to turn on the equivalent protocol for IPv6: “ipv6 cef”

JUNIPER

In JUNOS IPv6 is enabled on a per interface basis, the steps in the following slides show how to enable an interface by specifying the address. If the address part of the command is omitted the interface will enable itself with a link-local address only.

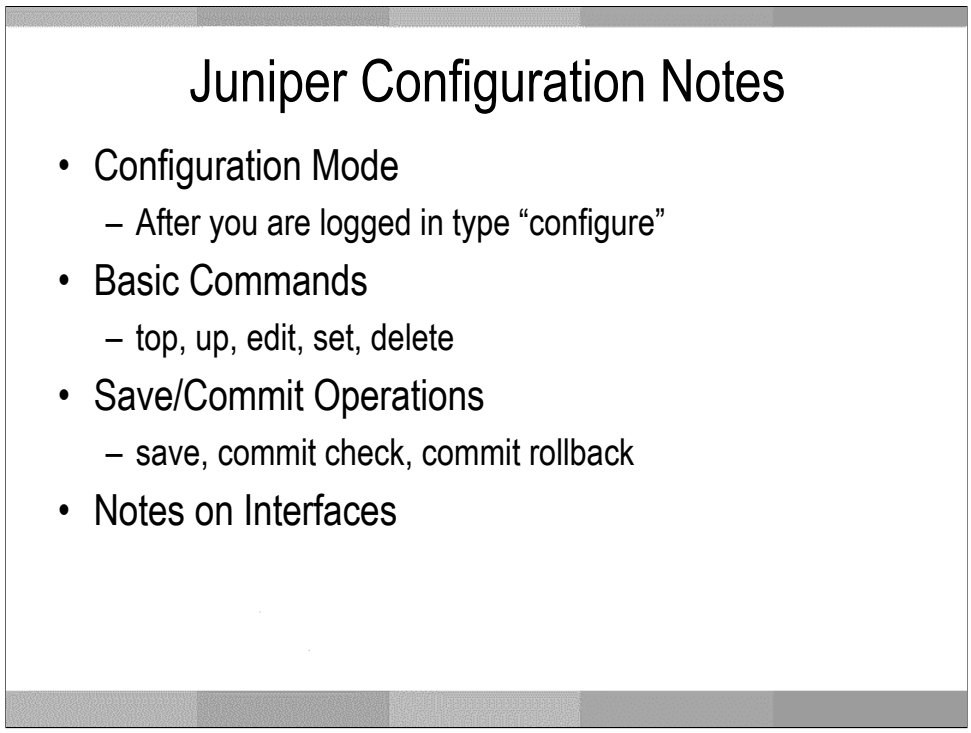

JUNOS operates a hierarchical system much like IOS and this can be moved around by using the “edit” command, e.g. “edit interfaces” will take you to the interfaces section of the configuration, “show” will now only display configuration applicable to this section. Further edit commands can be performed one after the other until reaching the lowest element in the tree. Edit commands can also be chained, e.g. “edit interfaces fe-1/3/0 unit 0”. Note that all interfaces have units even if they only have one unit, which is different from the way IOS works. During these labs we will be specifying which parts of the hierarchy that you need to be within before executing each command, for example:

From [edit interfaces fe-0/0/0]

“set unit 0 ……..”

Other commands such as “up” and “top” are also useful for browsing the hierarchy.

Saving/Commiting

Any saves or commits whilst in the hierarchy will only apply changes in that part of the tree, use “top” to return to the global edit space, JUNOS usefully tells you where you currently are in square brackets, [edit] being the global configuration space. To remove a configuration entry replace set with delete. To change anything in the current Juniper candidate configuration (the one you are editing) we need to use the “set” command.

Physical Interfaces

In JUNOS you should make sure you are editing an interface which physically exists in your router as it is possible under JUNOS to configure an interface before you plug the module in. Physical interfaces which exist can be viewed outside of configuration mode with the command “show interfaces terse”

Ex2b. Upstream Interfaces

- Add an address
  - See notes below
- Test your connectivity
  - Exit configuration mode and try pinging your teams downstream interface on the head router (from the edge router):

    ping 2001:630:81:4X0::1
- **X** is always your group number

Enabling IPv6 on an interface is like enabling IPv4; add an “ipv6 address” directive to the interface.

CISCO

From exec/enable mode type “configure terminal”

From global configuration terminal:

“interface GigabitEthernet0/1”

“ipv6 address 2001:630:81:7100::1/64”

Or just:

“interface GigabitEthernet0/1 ipv6 address 2001:630:81:7100::1/64”

If using the first method of browsing to a particular place in the hierarchy before applying changes, the use of “exit” can be used to return to global configuration mode or “end” drops you up one level. A final exit drops you out of configuration mode to a level-15 access console where changes can be saved with the “write memory” command. Note how the prompt changes to inform you of the current location.

JUNIPER

Enter “configure” mode

From [edit]

“edit interfaces fe-0/0/0 unit 0”

“set family inet6 address 2001:630:81:4X0::2/64”

“show” (only interface fe-0/0/0 unit 0 is shown)

BOTH

Refer to your topology diagram for the correct interfaces and address to use.

The upstream interface is FastEthernet-0/0/0. For example, Team A’s upstream interface should have address 2001:630:81:410::2/64.

Ex2c. Downstreams

- Add IPv6 addresses to downstream-facing interfaces
- Watch ethereal on clients for traffic
  - Router advertisements
  - Neighbour discovery packets
    - … particularly DAD
- Watch node interfaces for addresses
- ping6 using *globals* internal to subnet
  - Compare against using fe80::/10 link-locals

With the router demonstrating that it has native IPv6 connectivity, the next task is to extend that connectivity toward our client subnets. First, ensure that you have a traffic sniffer running on your client nodes and confirm that you have no global (i.e. non-fe80::/10) addresses configured on the nodes' fixed interfaces. (You could try running ifconfig/ipconfig to check this).

BOTH

In each of the downward sub-interfaces, add "ipv6/inet6 address" directives as per the network topology diagram for your group. Note that these are /64 subnets.

JUNIPER (extra step required)

JUNOS does not advertise a prefix to the clients until configured to. To do this:

From [edit]

"set protocols router-advertisement interface fe-1/0/X prefix 2001:630:81:4XX::/64"

"show protocols"

Note the lack of unit number here relating to the interface. There are other options available under protocols router-advertisement per interface such as default lifetime and current hop limit.

Observe the output of the traffic sniffer and then, in a different shell window, re-examine the nodes' interface configurations. (note it might take some time for the clients to receive the router advertisement).

At this point your nodes should have two IPv6 addresses of different scope: link-local and global. Try pinging and traceroute'ing between hosts on the two subnets of your team. There is currently no external route configured on the edge router so IPv6 nodes such as the other teams or sites on the internet (e.g. www.ipv6.ac.uk, www.kame.net) are not accessible, that will be the next step.

Ex2d: Testing your interfaces

- Below we have suggested a series of tests you can run at this stage to find what is now enabled
- Try and answer some of the questions below before you run the tests

Is it possible to ping using a link-local address a node using the destination's global address as target?

What about pinging a link-local address using a global address as source?

What do the packets look like on the wire as captured by the traffic filter?

What is the sequence of packets observed on the subnet if you try to ping the global address of a node under your subnet prefix that does not exist (e.g. If Team A client 2 tries to ping 2001:630:81:412::beef)?

When you try to ping an external host e.g. 2001:630:d0:f000::1 note the difference in response compared to pinging a node in your subnet that does not exist.

What does the IPv6 routing table look like on the client nodes? (hint: "ip route show" or "netstat -nr" on linux; "netsh interface ipv6 show routes" on Win32)

Ex2e. Static Routing

- Add a default route toward the head-router
- Do some more ping tests and trace(path/route) around the lab, to other groups, and beyond
- netstat -nr on clients, etc. to examine routes on clients

In its current state, the router is advertising a prefix to the two client subnets and has an interface over which it can reach the upstream router; however, the router knows nothing of how to route to destinations outside of these three connected subnets. In typical deployments, one might proceed with the use of an interior routing protocol such as IS-IS or OSPF. However, for the purposes of this introductory exploration into IPv6 today, we will install a single default route up to the head router.

The syntax is familiar, in the global router context simply add an ipv6 route for ::/0 to 2001:630:81:4X0::1, where 'X' is your group number corresponding to the point-to-point interface:

CISCO

```
ipv6 route ::/0 2001:630:81:4X0::1
```

JUNIPER

```
set routing-options rib inet6.0 static route ::/0 next-hop 2001:630:81:4X0::1
```

With this route installed, repeat the ping tests and interface/routing table observations from the previous element of the lab, but include some off-site destinations (e.g. www.6net.org, www.ist-ipv6.org)

Ex2e: IPv6 in Use

- You now have native IPv6 to your clients, experiment and see it at work!
- Leave Ethereal running to see the traffic
- DNS lookup, traceroute6, and browse to:
  - www.ipv6.org.uk/workshop/, www.uk6x.com, www.kame.net
- SSH into neighbouring machines, observe source addresses

DNS lookups under Windows can be done with “nslookup”. Run nslookup to get a prompt. At this prompt, type domain names to lookup. By default will find the A record (IPv4 address), can tweak this with:

> set type=aaaa

For AAAA records, or

> set type=all

To get both, if available.

Under Linux, use “host”. “host www.ecs.soton.ac.uk” will give the IPv4 address by default, use “host -t aaaa” to request AAAA (IPv6) records. Alternatively, “dig” is another tool to consider, particularly in trace-mode where the resolver forces an iterative walk of the DNS tree explicitly querying each authoritative server on the path toward the target label. For example, “dig +trace www.ist-ipv6.org aaaa” will use the locally-configured resolver (per /etc/resolv.conf) to first query the root, then .org SOAs, and so on. When the local network is dual-stack, occasionally one can observe IPv6 transport being used successfully to query a remote NS.

traceroute6 will show the routes that traffic will take. You can compare this to the IPv4 “traceroute” program.

Use a web browser (IE or Firefox under Windows should be fine, as should Firefox under Linux) to look at the web sites, listed on the slide. Your challenge for this lab is to be the first to visit http://www.ipv6.org.uk/workshop/ and see the waving Union Jack flag.

SSH should use IPv6 by default. Use “w” or “finger” when logged into another machine to see the source address being used. Choose a machine to access by its global (2001:630:d0:7XXX::) or link-local (fe80::) address. The Windows clients have “Putty” installed to use as an ssh client.

Summary

- Basic configuration and setup IPv6
- Router advertisements
- Neighbour discovery
- Static routes

- NEXT : IPv6 & Services