



JANET

IPv6 Hands-on Workshop

Module 5: Intradomain Routing (OSPFv3, IS-IS)

UKERNA, Lancaster University
and University of Southampton, 2006



IPv6 routing protocols

- IPv6 generally has the same routing protocols that IPv4 does today
 - RIPng, IS-IS, OSPFv3, BGP
 - We assume you have some familiarity with these protocols
 - If not, someone in your group should have 😊
- In the next two sessions we aim to look at:
 - Specifics for configuring IPv6 routing protocols
 - Considerations for routing protocols on dual-stack routers
- In the hands-on labs you will be able to try:
 - IS-IS (intradomain scenario)
 - BGP (interdomain scenario)

Intradomain protocols

- We begin with theory and hands-on for intradomain IPv6 routing
 - Protocols you may already run for IPv4 within your campus or your RNO backbone
 - The JANET core runs IS-IS for both IPv4 and IPv6 today
- We focus on the link-state protocols
 - OSPFv3
 - RFC2740
 - IS-IS
 - draft-ietf-isis-ipv6-06
- Both available on most router platforms
 - Including Cisco and Juniper

Intradomain routing protocols

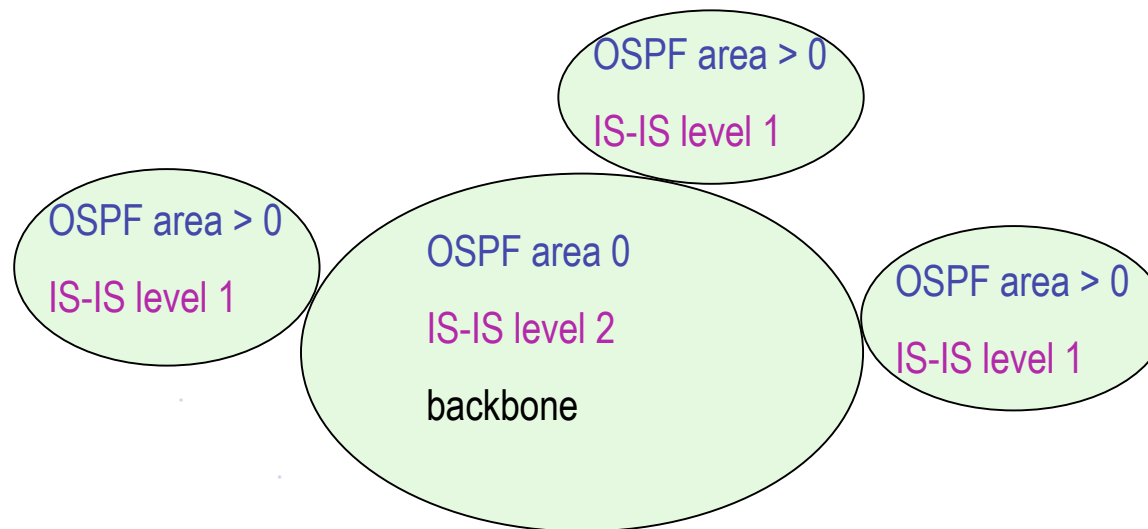
- Link State
 - OSPF, IS-IS
- Distance Vector
 - RIP, IGRP, EIGRP
- For IPv6: OSPFv3, IS-IS, RIPng and EIGRP
- Can mix and match routing protocols for IPv4 and IPv6
 - E.g. OSPF for IPv4 and IS-IS for IPv6
 - Would normally run two completely separate routing instances, ships in the night
 - With IS-IS it is possible to use a single common instance for both
 - Originally not possible to have separate IS-IS instances for IPv4 and IPv6, but Cisco IOS and JUNOS now have multi-topology IS-IS that allows this
- IPv6 configuration is similar to IPv4

OSPF/IS-IS

- Both are link state protocols
 - Faster convergence than distance vector protocols
 - Each router announces information on addresses and metrics for their neighbours
 - This information is flooded throughout a so-called area
 - Each router in the area has exactly the same information in a link state database
- Both use 2 level hierarchical model
 - Backbone (area 0 / level 2), and “non-backbone” areas. Information is aggregated at area boundaries. Topology of an area is invisible to routers outside
- Recommend a common IS-IS instance for both IPv4 and IPv6;
 - Or using OSPFv2 for IPv4 and OSPFv3 for IPv6
- Some networks have converted to IS-IS to have a single routing process. Less resources, easier to manage
 - Also a reason why the JANET core switched from OSPF to IS-IS

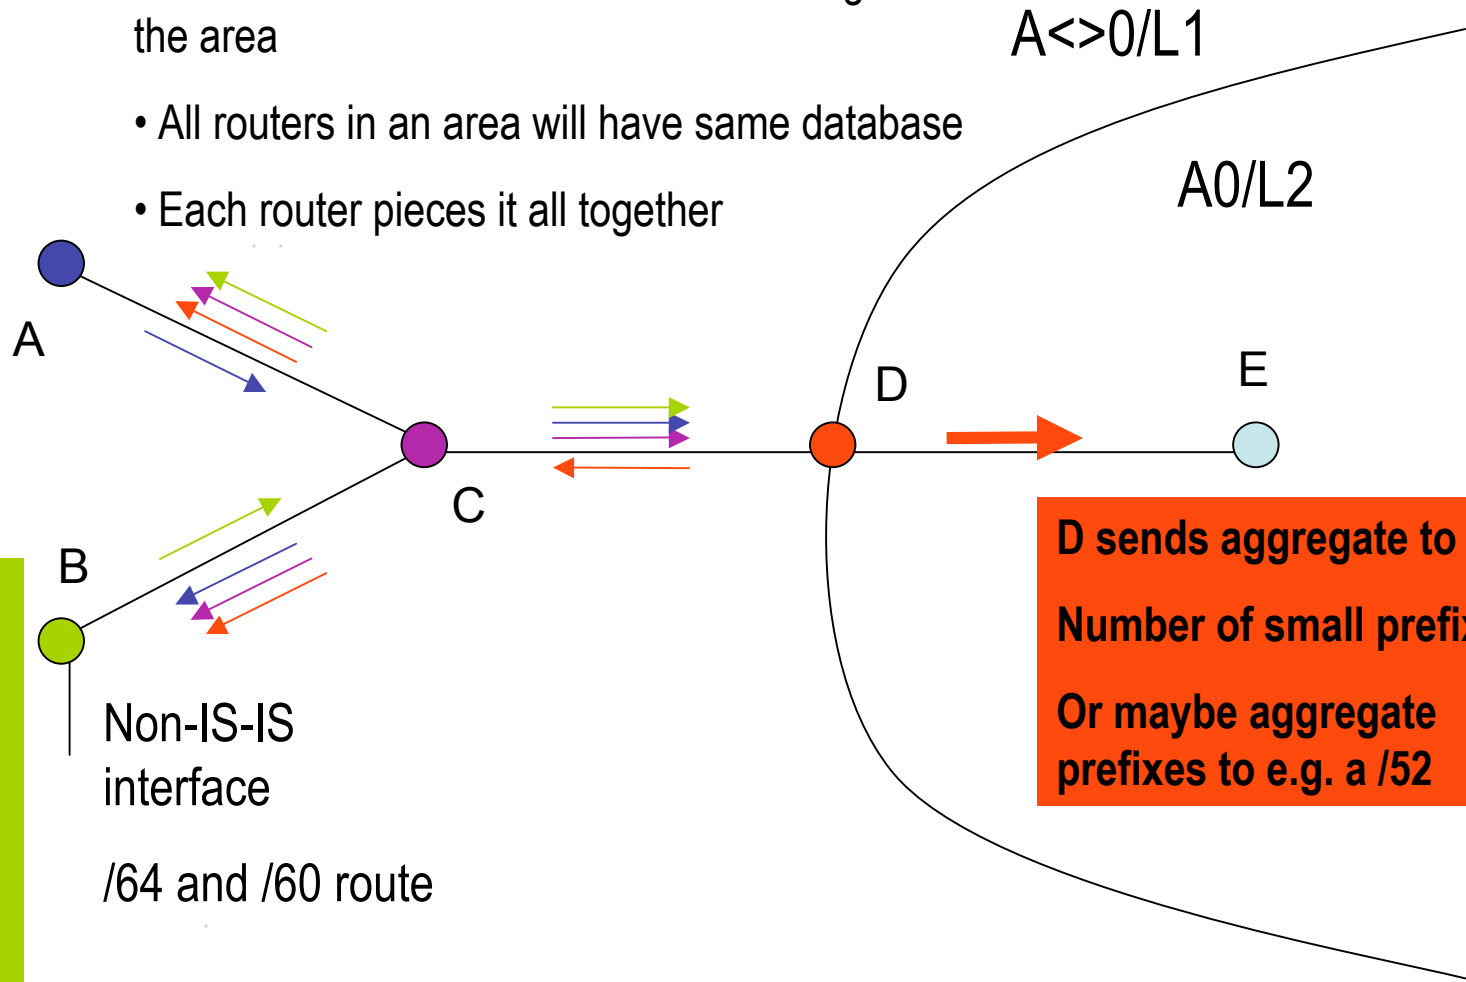
OSPF/IS-IS hierarchy

- Two level hierarchical architecture, may choose to have just one
 - Note that we are not considering how to use the two together, all areas use the same protocol
- The hierarchy is optional, with < 100 routers you are probably ok with just **area 0 / level 2**
 - Note that each router can also do aggregation for routes it announces
 - With more, it might be good to limit how much state is passed between the routers and how much state each router needs to cope with (in particular limiting state changes)
- Some example hierarchies
 - A very large university with a distributed campus could let each campus be **area > 0 / level 1** and a university core network **area 0 / level 2**
 - If there were no UK regional networks, but one management domain, then each region could be one area (**area > 0 / level 1**) and JANET core **area 0 / level 2**



How the hierarchy works

- Each routers local view is flooded throughout the area
- All routers in an area will have same database
- Each router pieces it all together



B's view:
Neighbour C
One /64
Another /64
The /60

Non-IS-IS
interface
/64 and /60 route

**D sends aggregate to E:
Number of small prefixes
Or maybe aggregate
prefixes to e.g. a /52**

OSPFv3

- OSPF for IPv6, some enhancements of OSPFv2
- Can run alongside OSPFv2 for IPv4
- Distributes IPv6 prefix reachability
- Makes use of IPv6, uses link-local addresses

- Two new Link State Advertisement (LSA) types
- Link LSA
 - Router's link-local address and prefixes for an interface
 - Only flooded on the link
- Inter-Area-Prefix LSA
 - Prefixes for stub and transit networks

Cisco IOS OSPFv3 configuration

- Must be enabled on all interfaces where it's wanted
 - `ipv6 ospf <process ID> area <area ID>`
 - By default off
 - That's all for simple setups
- May need some global config
 - `ipv6 router ospf <process ID>`
 - For edge interfaces
 - `passive-interface <interface name>`
 - Setting router-id, needed if no interfaces have IPv4 address
 - Defining area summary prefixes to have aggregation between areas
 - `area <area ID> range <prefix>/<prefix length>`

JUNOS OSPFv3 configuration

- Enabling on interfaces:
 - `set protocols ospf3 area <area id> interface <interface>`
 - By default off
 - That's all for simple setups
- No need to configure globally (check)
 - For edge interfaces
 - `set protocols ospf3 area <area id> interface <interface> passive`
- Aggregation of routes
 - `set protocols ospf3 export <policy statement name>`
(see IS-IS hands-on for details on policy statements)

Cisco IOS OSPFv3 status commands

- Showing neighbours:
 - `show ipv6 ospf neighbor`
 - `show ipv6 ospf [<process ID>]`
 - `clear ipv6 ospf [<process ID>]`
- Showing new LSA:
 - `show ipv6 ospf [<process ID>] database link`
 - `show ipv6 ospf [<process ID>] database prefix`
- Debug commands:
 - `debug ipv6 ospf ...`

JUNOS OSPFv3 status commands

- Showing neighbours:
 - `show ospf neighbor (brief|detail...)`
- Showing new LSA:
 - `show ospf3 database (...link|network|summary...)`
- Debugging
 - `show ospf3 log (...)`
- Other useful commands:
 - `show ospf3 interface (...)`
 - `show ospf3 route (...)`
 - `show ospf3 statistics (...)`

OSPF References

- For further details on OSPF configuration on IOS:
 - http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a00801d660d.html
- On JUNOS:
 - <http://www.juniper.net/techpubs/software/junos/junos80/swconfig80-routing/download/ospf-config.pdf>
 - or
 - <http://www.juniper.net/techpubs/software/junos/junos80/swconfig80-routing/frameset.htm>
- Also see the IETF RFCs
 - RFC 2740 – OSPF for IPv6 (OSPFv3)
 - RFC 2328 – OSPFv2

IS-IS

- RFC 1195 specifies IS-IS for IP
 - It was made so that one process could take care of both IP and CLNS
- IS-IS for IPv6 is defined
 - See draft-ietf-isis-ipv6-06
- Advantage is you can have one single process for IPv4 and IPv6 routing
- Language is different but concepts mostly the same as OSPF
 - IS-IS level 2 = OSPF area 0. IS-IS level 1 is “not backbone” area
 - IS-IS adjacency = OSPF neighbour
 - IS-IS DIS = OSPF DR

IS-IS and OSI

- In OSI speak hosts are End Systems (ES) and routers are Intermediate Systems (IS)
 - So IS-IS is a router to router protocol
- ISO 10589 specifies IS-IS for CLNS
 - Uses TLV coding, easy to extend
- Don't need to know about OSI stuff to run IS-IS
 - But will need to construct OSI NSAP addresses
 - More on this later...

IS-IS for IPv6

- Introduces two new TLVs for IPv6 routing
- IPv6 Reachability TLV
 - External bit
 - Equivalent to IP Internal/External Reachability TLV's
- IPv6 Interface Address TLV
 - For Hello PDUs, must contain the Link-Local address
 - For LSP, must only contain the non-Link Local address
- IPv6 NLPID is advertised by IPv6 enabled routers
 - Network Layer Protocol ID

Single SPF is nice, but...

- Single SPF per level for OSI, IPv4 and IPv6
- Interfaces configured with IS-ISv6 must support IPv6
- Interfaces configured with IS-ISv4 must support IPv4
- All routers in an area supposed to have same protocols
 - All IPv4-only, all IPv6-only, or all with both
- By turning off adjacency check one can lessen this restriction
 - e.g. if all IPv4, then set of routers also having IPv6 must be convex
 - i.e. shortest path between two IPv6 routers must not be via IPv4-only router
- Also, two areas with IPv6 routers should not be connected by area with IPv4-only (again, convex...)
 - Note that IS-IS does not prohibit this (even with adjacency check on), but you will get a black hole

Constructing NSAP addresses

- When using IS-IS for IPv4/IPv6, you only see OSI's NSAP addresses as router and area ID's
- Can be constructed in several ways
- We suggest the private coding below, see RFC 1195 for other alternatives

49 .xxxx .yyyy .iiii .iiii .iiii .00

- 49 is AFI (Authority and Format Indicator), 49 means private
 - So we can do what we want really
- Next we have 1-12 bytes for domain and area
 - Suggest AS number in hex (the x's)
 - Followed by 4 bytes for area ID (the y's). So each area different id
- Then there must be 6 bytes System Identifier (the i's)
 - Can be MAC address, IPv4 address in BCD, or part of IPv6 address, ...
 - This must be unique within the routing domain (within AS)
- Finally we have 00 which is sort of like IP port numbers. Should be 0
- Below is an example where AS number is 64646 (fc86 in hex), area ID is 1 and routers IPv4 loopback address is 158.38.0.128

49 .fc86 .0001 .1580 .3800 .0128 .00

IOS IS-IS configuration

- First you need to configure the IS-IS router process
- Using NSAP address constructed per previous slide, we have

```
router isis  
net 49.fc86.0001.1580.3800.0128.00
```

- IS-IS must be explicitly enabled on interfaces
- You would usually enable it only on interfaces facing other IS-IS routers
 - You should not enable it at the edge where there are only hosts
 - You would normally enable both IPv4 and IPv6 IS-IS on the same interfaces

```
ip router isis  
ipv6 router isis
```

JUNOS IS-IS Configuration

- Specifying the router with NSAP address:

```
set interface lo0 unit 0 family iso address <NSAP addr>
```

- Enabling IS-IS on interfaces:

```
set interface <interface> family iso
```

Example IOS IS-IS configuration

```
router isis
```

```
net 47.0158.0038.0000.0128.00
```

```
is-type level-2-only
```

```
log-adjacency-changes
```

```
summary-address 158.38.0.0 255.255.0.0
```

```
...
```

```
address-family ipv6
```

```
no adjacency-check
```

```
redistribute connected
```

```
redistribute static
```

```
default-information originate
```

```
summary-prefix 2001:700:0:500::/56
```

```
exit-address-family
```

- is-type can be **level-1**, **level-1-2** or **level-2-only**
- If need only one area, use **level-2-only**, default is 1-2

Summary-address/prefix can be used for aggregation

IPv6 specific stuff inside **address-family ipv6**

Turn off check only if not all neighbours do IPv6

Here we redistribute prefixes for connected interfaces and our static routes

Tells router to originate a default-route

Example JUNOS IS-IS configuration

```
Protocols {
  isis {
    level 1 disable
    export [ redistribute-connected aggregate-
            addresses]
    no-ipv4-routing
    interface lo0
    interface ...
  }
}

Interfaces {
  fe-1/0/0 {
    family iso;
  }
}
```

Use level 2 only, default is 1-2

Policies applied to this interface. (see lab)

Spceifices interface supporting IS-IS

Enables an interface.

Checking IS-IS status in IOS

- Look for neighbours:
 - **sh clns neigh**
 - **sh isis neigh**
- Check topology learnt by router:
 - **sh isis top**
- Show full link-state database:
 - **sh isis database [<verbose>]**
 - **sh isis database detail**
- Debugging:
 - **debug isis ...**

Checking IS-IS status in JUNOS

- Checking topology learnt by router:
 - `show isis route`
- Show full link-state database:
 - `show isis database`
- Other useful commands:
 - `show isis adjacency`
 - `Show isis statistics`
 - `Show isis hostname`
 - `Show isis interface`

Multi-Topology IS-IS

- If you can't fulfill the single SPF constraints
 - e.g. you don't have convex sets of IPv6 routers
- With multi-topology IS-IS you can have separate IS-IS instances for IPv4 and IPv6
- Introduces several new TLVs
- Some complexities if start with this and later want to move to a single topology
- We recommend a single topology where possible
 - Current practice on the JANET core

IS-IS References

- For further details on IS-IS and configuration on IOS:
 - http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a00801d65f6.html
- For IS-IS on JUNOS see:
 - <http://www.juniper.net/techpubs/software/junos/junos80/swconfig80-routing/download/isis-config.pdf>
- JUNOS and IOS paper:
 - http://www.juniper.net/solutions/literature/app_note/350005.pdf
- RFCs:
 - RFC 1195 – Use of OSI IS-IS for Routing in TCP/IP and Dual Environments
 - IETF draft-ietf-isis-ipv6-06 – Routing IPv6 with IS-IS