COMP312-09A
Communications and Systems Software

Routing 2 – IGPs

Richard Nelson

richardn@cs.waikato.ac.nz

Static Routing

- Static Routing (i.e. manual) is commonly used in IP networks.
- Networks that have no redundancy may not benefit from automatic route changes.
- Static routes provide the network operator complete control of the routing and are easier to understand.
- Static routes are not vulnerable to signalling loss or software bugs and are more likely to restart after a software or hardware failure.
- Complex addressing schemes may fool some routing protocols (especially older class-based protocols)
- Static routes require no signalling may be preferred between routers managed by different organisations.
- Static routes have no metric so can be used to load balance over arbitrary routes.

Dynamic Routing

- IP Routing protocols provide Dynamic Routing:
- Often requires less configuration and re-configuration.
- Automatically reconfigures as the network changes topology.
- Can generally be controlled through careful selection of the link metrics.
- Modern routing protocols are dynamic in terms of response to topology changes, not in response to traffic flows.

Routing Protocols

- Interior Gateway Protocol (IGP)
  - Used inside an organisation's network (except a network so small it can use static routes).
  - These protocols assume that all routers are entitled to any knowledge about the internal topology of the network.
  - They aim to optimise the network.
- Exterior Gateway Protocol (EGP)
  - Used between organisations. EGP routing is influenced less by network topology and more by business decisions.

Interior Gateway Protocols

- Distance Vector
  - RIP v2 – Still used in very small networks.
  - EIGRP – Used in some corporate networks. Cisco proprietary.
  - Many older and proprietary protocols.
- Link State
  - OSPF – Usual Interior Gateway Protocol for all but very small or very large networks.
  - IS-IS – Used by some telcos

Exterior Gateway Protocol

- **BGP**

IGPs

- Default Routes
- Distance Vector Protocols
  - RIP
- Link State Protocols
  - OSPF

Default Routes

- A default route is the one to be used if no other route to that destination is known.
- Usually written 0.0.0.0/0
- Nearly all computers are statically configured with a default route or “gateway” address. This is generally the address of a reachable interface on a router which has connectivity to the outside.

Default Route - Example

```
[dmneal@zen ~]$ netstat -nr
Kernel IP routing table
Destination     Gateway         Genmask         Flags   MSS Window  irtt Iface
192.168.1.0     0.0.0.0         255.255.255.248 U         0 0          0 eth0
169.254.0.0     0.0.0.0         255.255.0.0     U         0 0          0 eth0
0.0.0.0         192.168.1.1     0.0.0.0         UG        0 0          0 eth0
```

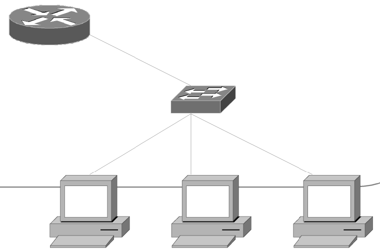

Default Routes

- Computers (end hosts) almost never run routing protocols.
- Routers commonly also have a default route, which
  - May be statically configured
  - May be advertised from outside using a routing protocol
  - May be chosen by the router from among more than one possible default routes

IGPs

- Default Routes
- Distance Vector Protocols
  - RIP
- Link State Protocols
  - OSPF

Distance Vector IGPs

- Routers exchange their view of the rest of the network with their neighbours.
- Each router chooses the next hop based on the network views of its neighbours.
- Sometimes called routing by rumour.
- All Distance Vector protocols use techniques to mitigate the counting to infinity problem and slow convergence time of the basic Bellman-Ford algorithm.

Routing Information Protocol -RIP

- Was the culmination of early routing experience with the ARPANet.
- Standardised by the routed program in the 4.2BSD released in 1982.
- Documented in RFC 1058 in 1988.
- Is a class based protocol and uses simple hop count metric (∞ = 16).
- Was very popular due to widespread free distribution.
- Now obselete

Routing Information Protocol v2

- Distance Vector IGP, defined in RFC2453 (obseletes RFC1723 from 1994)
- Classless version of RIP
- Comparatively simple to implement.
- Runs on very small, cheap routers.
- Does not scale beyond small networks.
- Fairly common in small branch offices.

Routing Information Protocol v2

- A router running RIP sends out all of its RIP routes in a "response" packet.
- A RIP router may on starting up generate a "request" message. A "response" message may be sent in reply to a request, but most are not.
- Response packets are sent when a route changes, or at fixed intervals, usually 30s.
- Response packets are multicast using UDP to address 224.0.0.9.

RIP Response Packets

A RIP response packet contains:

- A version number = 2
- May contain authentication information (plain text password or MD5)
- Route information:
  - Route tag – internal (RIP) or external (learned from elsewhere) route
  - IP address e.g. 192.168.1.0
  - Netmask e.g. 255.255.255.248
  - Next Hop
  - Metric

RIP Route Information

Route information:

- Route tag – internal (RIP) or external (learned from elsewhere) route
- IP address e.g. 192.168.1.0
- Netmask e.g. 255.255.255.248
- Next Hop - IP address to which packets bound for this address range should be forwarded.
- Metric- Hop count. Valid values are 1 to 15, with 16 meaning unreachable.

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

RIP in Small Offices

RIP runs in very small, cheap routers, so is still to be found in small offices.

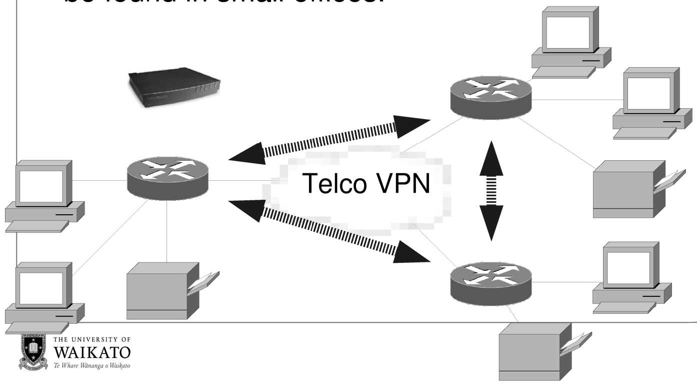

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

EIGRP

- Enhanced Interior Gateway Routing Protocol.
- Cisco classless replacement for IGRP.
  - Itself an improved version of RIP
- Uses advanced calculation technique that is loop-free and fast converging.
- Improved security.
- Multiprotocol (IPX, Appletalk).
- Used in many all-Cisco enterprise networks in the 1990's.

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Distance-Vector IGPs - Summary

- The original basis for ARPANet routing.
- Still very common in simple networks due to the simplicity and low cost of RIP.
- Computationally lightweight but Bellman-Ford algorithm has practical problems.
- Only proprietry versions evolved to be suitable for large networks.

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

IGPs

- Default Routes
- Distance Vector Protocols
  - RIP
- Link State Protocols
  - OSPF

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Link State IGPs

- Routers share information about their local links with all other routers in the network.
- Each router builds a link state database and runs Dijkstra's algorithm to calculate shortest paths.
- Three distinct phases
  1. Neighbour discovery
  2. Link-State Advertising
  3. Shortest path calculation

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Link-State Advertising

- Link state information needs to be spread to all other routers when the paths are not necessarily known or accurate.
- Selective Flooding is used to reach all destinations irrespective of the state of the network knowledge.
- Sequence numbers are used to limit flood traffic and maintain current information.

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Link-State Advantages

- There are two main advantages to Link-State protocols.

  1. Link state advertisements do not need to be processed before being passed on so topology change information propogates as quickly as possible to the entire network.

  2. Each router has a full topology database and so can chose loop-free routes.
- The relative signalling costs depend on the network. Link state protocols send packets to the whole network, but packets are smaller and do not require multiple exchanges for network convergence.
- The complete topology database may also have advantages if implementing advanced network services.

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Link State Disadvantages

1. Complexity
2. Processing Overhead

- Large networks are increasingly complex anyway and modern CPUs have little difficulty with the running SPF over large topologies so these disadvantages are seen as minor.

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Link-State Protocols

- There are two Link-State routing protocols used for IP.
  1. OSPF
  2. IS-IS
- Link-state routing has also been used for Asynchronous Transfer Mode (ATM) networks.

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Open Shortest Path First

- Main IETF standardised (and recommended) IGP.
- Main version (2) first specified in RFC 1247 (1991) and has been updated several times, most recently in RFC 2328 (1998).
- Version 1 was experimental, Version 3 adds IPv6 support
- Dedicated IP classless routing protocol.
- Metric is bandwidth or (in effect) hop count.
- Can be scaled further using areas (though this is not common).
- Can use more than one route to the same destination.
- Usual IGP in use today.

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

Open Shortest Path First

- Each OSPF router forms neighbour relationships - “adjacencies” - with the other OSPF routers around it.
- Once an adjacency has been brought up, topology (not route) information is exchanged. Mostly.
- After each router's topology database has been populated, it uses the SPF (Dijkstra) algorithm to calculate the routes it will use.
  - SPF = Dijkstra – Routing Algorithm
  - OSPF – Routing Protocol

THE UNIVERSITY OF WAIKATO
Te Whare Wānanga o Waikato

OSPF Phases

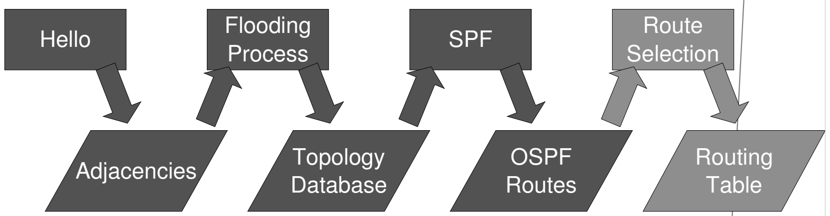

OSPF and IP

- OSPF v2 is defined around IPv4, so new protocol version, OSPF v3, required for IPv6. If you run both IP versions, you need both OSPF versions.
- Does not use either UDP or TCP. OSPF is a protocol defined to run immediately over IP.

Forming an Adjacency

- The router is configured to run OSPF on some or all interfaces.
- Out of all configured interfaces, a Hello packet is multicast to 224.0.0.5 (AllRouters). This includes
  - Its router ID (usually a loopback interface address)
  - Flags describing what OSPF capabilities the sender has
  - An OSPF priority
  - A list of addresses this router has seen Hello packets from.

Forming an Adjacency

- The adjacency state “Full” indicates that this process has completed.
- Adjacencies are kept alive by the receipt of Hello (“keepalive”) packets.
- Each router uses it's adjacencies to form a description of it's local topology called a *Link State Advertisement* (LSA).

Forming an Adjacency

- Once routers have received Hello packets containing their own ID's, Data Description packets are exchanged containing descriptions of each router's topology database.
- The topology database is made up of Link State Advertisements (LSA's).
- If a router finds its neighbour has an LSA it doesn't, or has a more recent version of the same LSA, it sends a Link-State Request packet and receives back a Link-State Update packet containing that LSA.

Designated Routers

- On a broadcast segment (e.g. Ethernet) which has more than one OSPF router on it, the router with the highest priority will be elected as the Designated Router (DR). The router with the second highest priority becomes the Backup Designated Router (BDR).
- In such a segment, routers form adjacencies only with the DR and BDR, cutting down the replication of information.
- Packets to be sent to the DR and BDR go to 224.0.0.6 (AllDRRouters).
- Packets from the DR to routers on the segment go to 224.0.0.5 (AllRouters).

Designated Routers

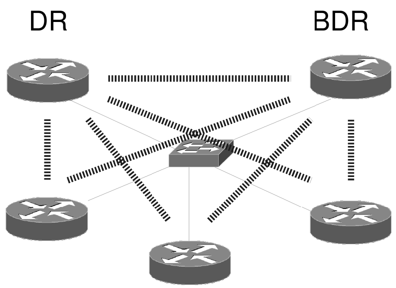

Link State Advertisements - Header

LS Age	
Options	LS Type
Link State ID	
Advertising Router	
LS Sequence Number	
LS Checksum	
Length	

20 Bytes total

LSA Types

In the simplest OSPF network, the LSA types you will see are:

1 Router LSA – From any router. Describes networks connected to the router

2 Network LSA – From a DR. Describes routers attached to that network

THE UNIVERSITY OF WAIKATO Te Whare Wānanga o Waikato

LSA Metrics

- Each LSA contains a metric for the link it describes.
- In protocol terms, this is just a 16-bit number (max. 65535).
- Could just use the same number for all links, so metric = hop count.
- Cisco and Juniper default value is 10^8/interface bandwidth, so 10 for 10Mbps ethernet, 1 for 100Mbps ethernet or faster. Or reconfigure the reference bandwidth to be a larger number.
- Each LSA has an age recorded. LSA's are periodically re-sent, typically every half hour. If a new version of the LSA is not received when the age reaches a pre-set limit, the LSA is deleted.

THE UNIVERSITY OF WAIKATO Te Whare Wānanga o Waikato

LSA Aging

- Each LSA has an age recorded and a sequence (i.e. version) number.
- LSA's are periodically re-sent, typically every half hour.
- If a new version of the LSA is not received when the age reaches a pre-set limit, the LSA is deleted.

THE UNIVERSITY OF WAIKATO Te Whare Wānanga o Waikato

OSPF Areas

- OSPF can be scaled up by introducing an element of heirarchy. This is done using areas.

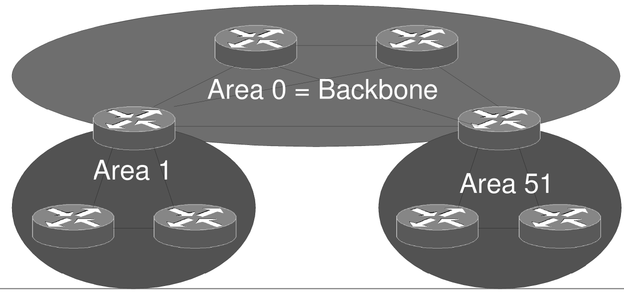

THE UNIVERSITY OF WAIKATO Te Whare Wānanga o Waikato

OSPF Areas

- The backbone area is always numbered 0.
- All other areas must be connected to the backbone area.

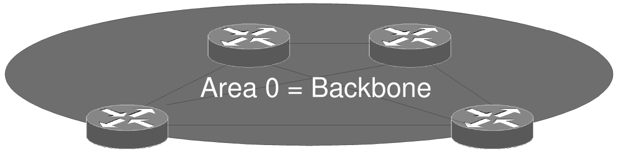

OSPF Areas

- An Area Border Router (ABR) has interfaces in more than one area. If addresses within a non-backbone area are taken from a specific address range (say, 130.217.16/20), only a single prefix need be advertised to the backbone by the ABR.
- Yes, that is a shift away from pure link-state routing.
- All ABRs have one interface in Area 0

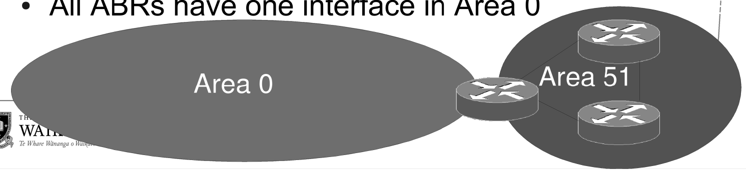

ASBR's and External Routes

- An Autonomous System Border Router may advertise routes learned outside OSPF.
- This is done using External LSA's.
- One use of this ability is for the router(s) connected to the outside world to advertise a default route.

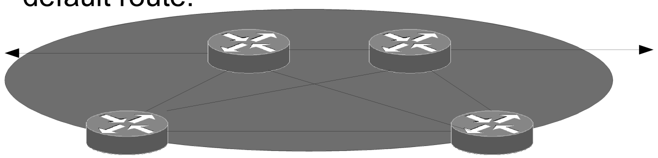

ASBR's and External Routes

If two ASBR's both advertise a default route with the same metric, what will happen?

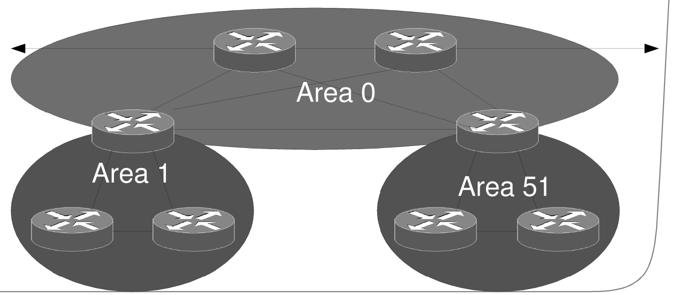

LSA Types

In a more complex OSPF network, you may see:

1 Router LSA – From any router. Describes networks connected to the router. Only flooded within an area.

2 Network LSA – From a DR. Describes routers attached to that network. Only flooded within an area.

3 Network Summary LSA – From an ABR. Describes an address range available in another area.

4 ASBR Summary LSA – From an ABR. Describes where an ASBR is.

5 External LSA – From an ASBR. Describes routes from outside this OSPF network. Flooded to all areas (except those configured not to receive external routes).

Further LSA types are not covered in this course.

Reading

Huitema, Routing in the Internet, Prentice Hall, 1995 pp. 65-90

Soricelli, JNCIA: Juniper Networks Certified Internet Associate Study Guide, Sybex, 2002 pp. 13-43 and Chapter 6 (See http://www.juniper.net/training/certification/books.html)

Moy, J.T. OSPF: Anatomy of an Internet Routing Protocol. Addison Wesley 1998

THE UNIVERSITY OF WAIKATO Te Whare Wānanga o Waikato