

Introduction to OSPF

ISP/IXP Workshops

Agenda

Cisco.com

- **OSPF Primer**
- **OSPF in Service Provider Networks**
- **OSPF BCP – Adding Networks**
- **OSPF Command Summary**

OSPF Primer

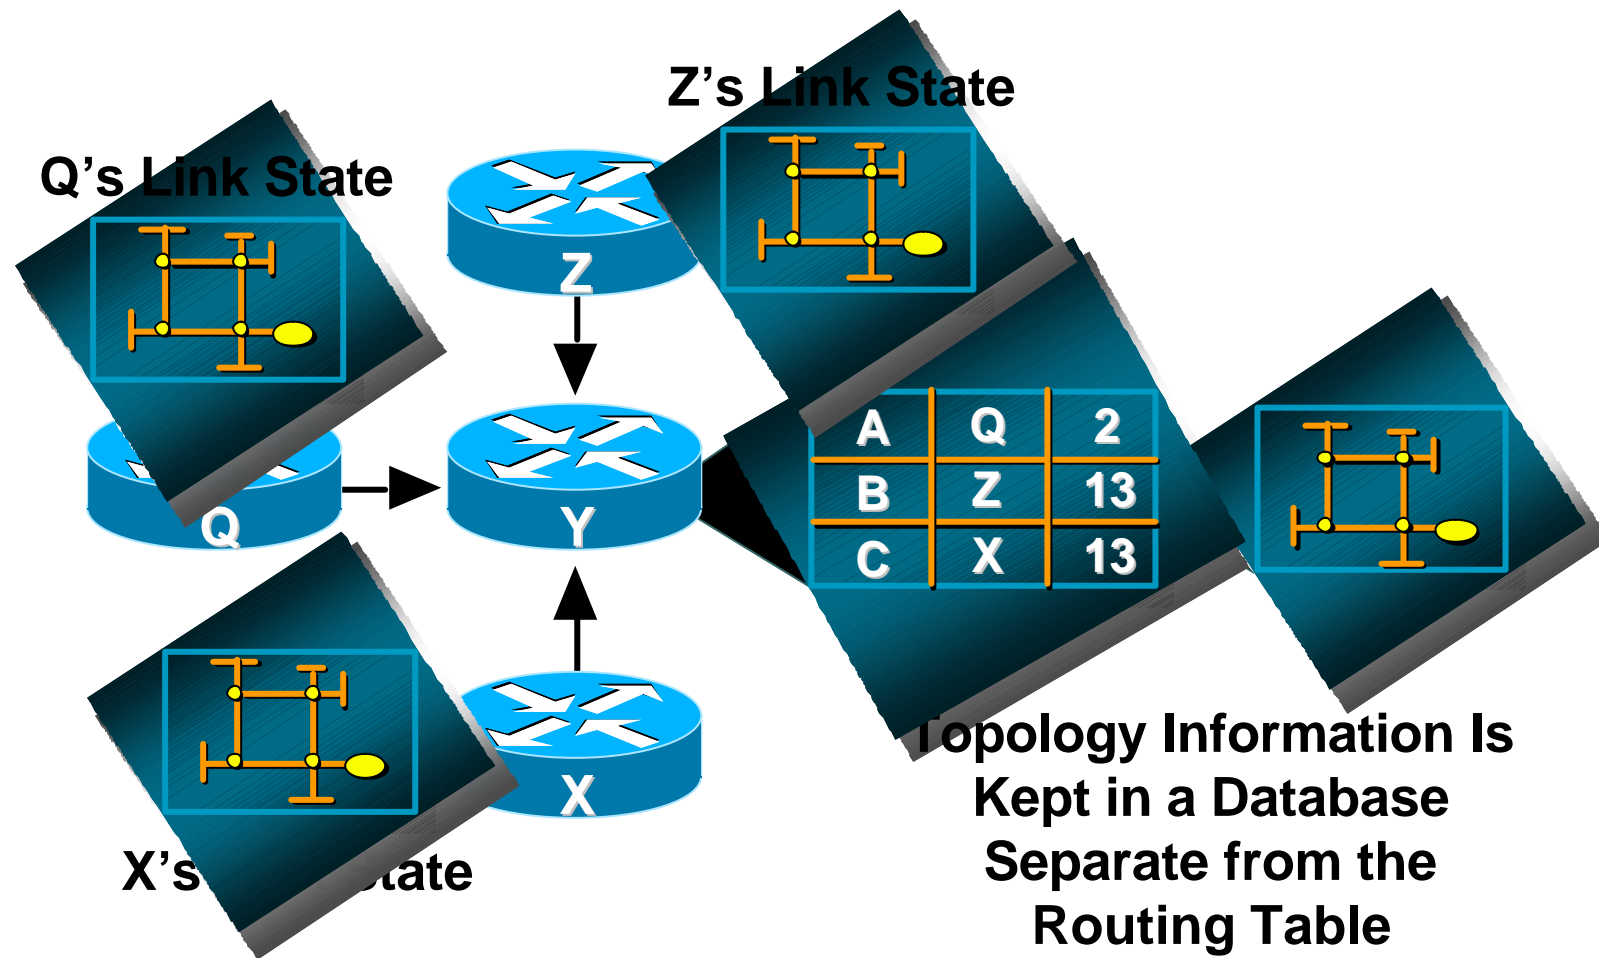
OSPF

Cisco.com

- **Open Shortest Path First**
- **Link state or SPF technology**
- **Developed by OSPF working group of IETF (RFC 1247)**
- **Designed for TCP/IP Internet environment**
- **Fast convergence**
- **Variable-length subnet masks**
- **Discontiguous subnets**
- **No periodic updates**
- **Route authentication**
- **Delivered two years after IGRP**
- **OSPF standard described in RFC2328**

Link State

Cisco.com




Link State Routing

Cisco.com

- **Neighbour discovery**
- **Constructing a Link State Packet (LSP)**
- **Distribute the LSP**
(Link State Announcement – LSA)
- **Compute routes**
- **On network failure**
New LSPs flooded
All routers recompute routing tables

Low Bandwidth Utilisation

Cisco.com



- Only changes propagated
- Multicast on multi-access broadcast networks


Cisco.com

The diagram illustrates a network topology for a distance vector protocol. It features four routers, R1, R2, R3, and R4, represented as blue circles with white arrows indicating internal routing. R1 is connected to N1 and N2. R2 is connected to R1 and R3. R3 is connected to R2 and R4. R4 is connected to R3 and N5. N2 and N3 are labeled "FDDI Dual Ring" and "Cost = 1". N4 is labeled "Cost = 10". N5 is labeled "Cost = 10".

Fast Convergence

Cisco.com

- **Detection Plus LSA/SPF**



Fast Convergence

Cisco.com


- **Finding a new route**

LSA flooded throughout area

Acknowledgement based

Topology database synchronised

**Each router derives routing table
to destination networks**



Utilises IP Multicast for Sending/Receiving Updates


Cisco.com

- **Broadcast networks**
 - All routers must accept packets sent to AllSPFRouters (224.0.0.5)**
 - All DR and BDR routers must accept packets sent to AllIDRouters (224.0.0.6)**
- **Hello packets sent to AllSPFRouters (Unicast on point-to-point and virtual links)**

OSPF Areas


Cisco.com

- **Group of contiguous hosts and networks**
- **Per area topological database**
 - Invisible outside the area
 - Reduction in routing traffic
- **Backbone area contiguous**
 - All other areas must be connected to the backbone
- **Virtual Links**



Classification of Routers

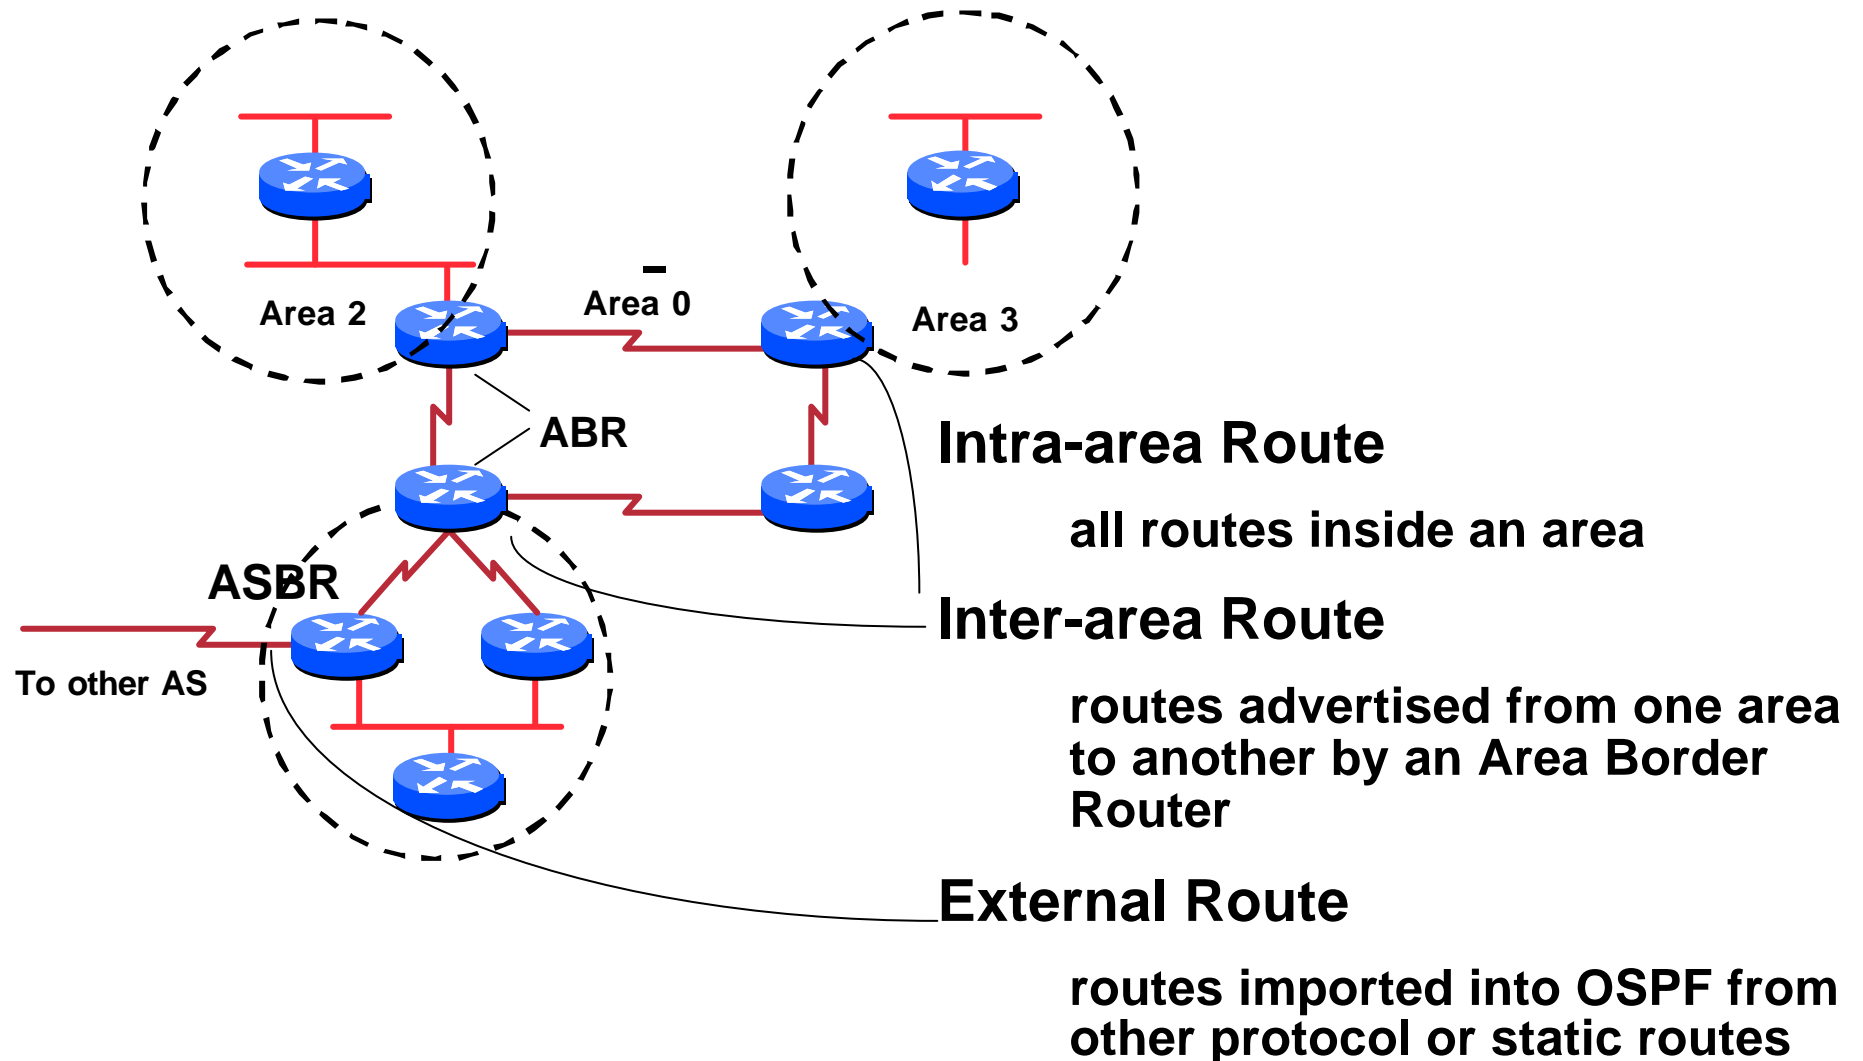
Cisco.com



- Internal Router (IR)
- Area Border Router (ABR)
- Backbone Router (BR)
- Autonomous System Border Router (ASBR)

OSPF Route Types

Cisco.com



Inter-Area Route Summarisation

Cisco.com


- Prefix or all subnets
- Prefix or all networks
- 'Area range' command

With
summarisation

Network	Next Hop
1	R1

Without
summarisation

Network	Next Hop
1.A	R1
1.B	R1
1.C	R1




External Routes

Cisco.com

- **Redistributed into OSPF**
- **Flooded unaltered throughout the AS**
- **OSPF supports two types of external metrics**

Type 1 external metrics

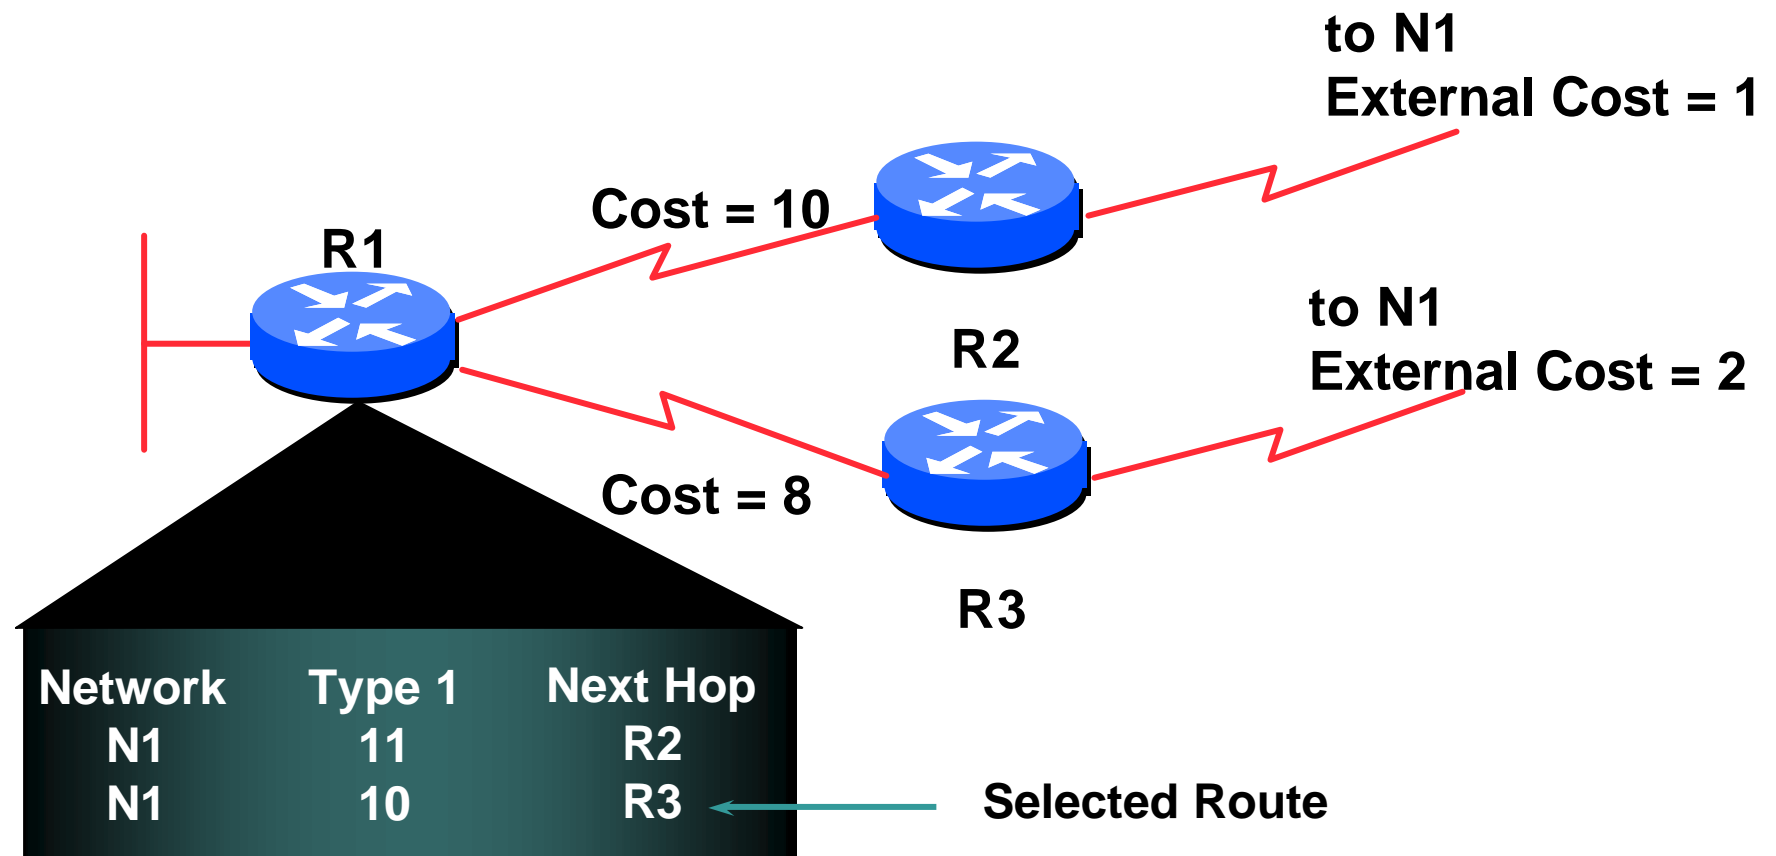
Type 2 external metrics (Default)



External Routes

Cisco.com

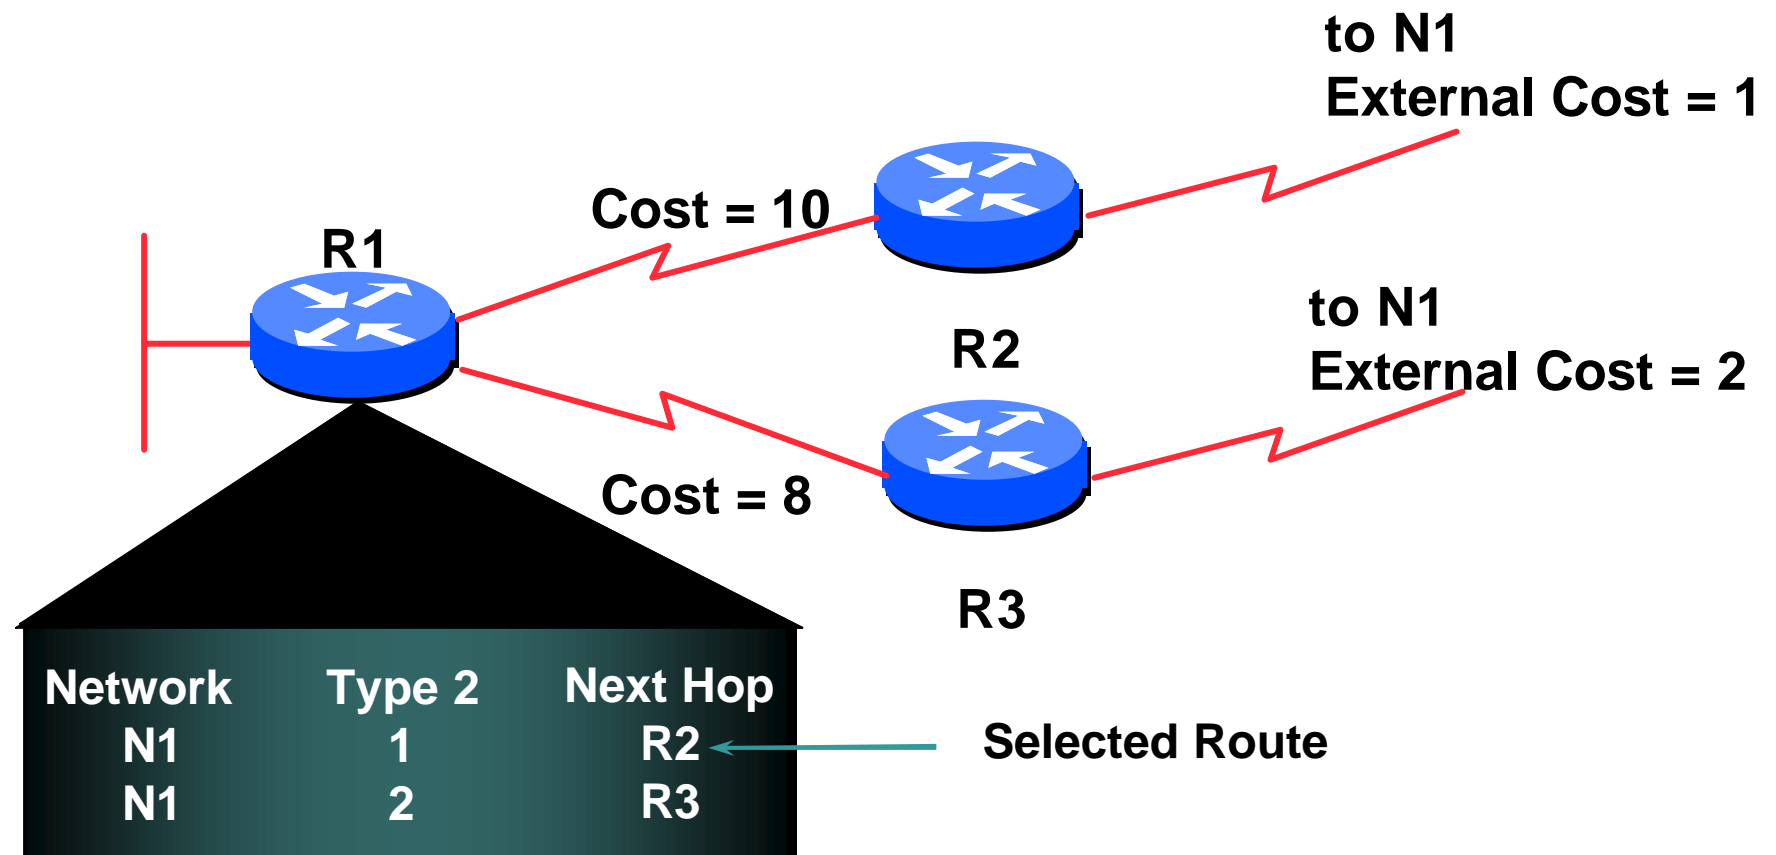
- **Type 1 external metric: metrics are added to the summarised internal link cost**



External Routes

Cisco.com

- **Type 2 external metric: metrics are compared without adding to the internal link cost**



Topology/Link State Database

Cisco.com

- **A router has a separate LS database for each area to which it belongs**
- **All routers belonging to the same area have identical database**
- **SPF calculation is performed separately for each area**
- **LSA flooding is bounded by area**


Protocol Functionality

Cisco.com

- **Bringing up adjacencies**
- **LSA types**
- **Area classification**

The Hello Protocol


- Responsible for establishing and maintaining neighbour relationships
- Elects designated router on multi-access networks



The Hello Packet

Cisco.com

- Router priority
- Hello interval
- Router dead interval
- Network mask
- Options: T-bit, E-bit
- List of neighbours




Designated Router

Cisco.com

- One per multi-access network

Generates network links advertisements

Assists in database synchronization




Designated Router by Priority

Cisco.com

- Configured priority (per interface)
- Else determined by highest router ID

Router ID is the loopback interface address, if configured, otherwise the highest IP address




Neighbouring States

Cisco.com

- **2-way**

Router sees itself in other Hello packets

DR selected from neighbours in state 2-way or greater



Neighbouring States


Cisco.com

- **Full**

Routers are fully adjacent

Databases synchronised

Relationship to DR and BDR




When to Become Adjacent

Cisco.com

- **Underlying network is point to point**
- **Underlying network type is virtual link**
- **The router itself is the designated router**
- **The router itself is the backup designated router**
- **The neighbouring router is the designated router**
- **The neighbouring router is the backup designated router**

LSAs Propagate Along Adjacencies

Cisco.com



- **LSAs acknowledged along adjacencies**

Routing Protocol Packets

Cisco.com

- **Share a common protocol header**
- **Routing protocol packets are sent with type of service (TOS) of 0**
- **Five types of OSPF routing protocol packets**
 - Hello – packet type 1**
 - Database description – packet type 2**
 - Link-state request – packet type 3**
 - Link-state update – packet type 4**
 - Link-state acknowledgement – packet type 5**

Different Types of LSAs

Cisco.com

- **Five distinct type of LSAs**

Type 1 : Router LSA

Type 2 : Network LSA

Type 3 and 4: Summary LSA

Type 5 and 7: External LSA

Router LSA (Type 1)

Cisco.com

- **Describes the state and cost of the router's links to the area**
- **All of the router's links in an area must be described in a single LSA**
- **Flooded throughout the particular area and no more**
- **Router indicates whether it is an ASBR, ABR, or end point of virtual link**

Network LSA (Type 2)

Cisco.com

- **Generated for every transit broadcast and NBMA network**
- **Describes all the routers attached to the network**
- **Only the designated router originates this LSA**
- **Flooded throughout the area and no more**

Summary LSA (Type 3 and 4)

Cisco.com

- **Describes the destination outside the area but still in the AS**
- **Flooded throughout a single area**
- **Originated by an ABR**
- **Only intra-area routes are advertised into the backbone**
- **Type 4 is the information about the ASBR**

External LSA (Type 5)

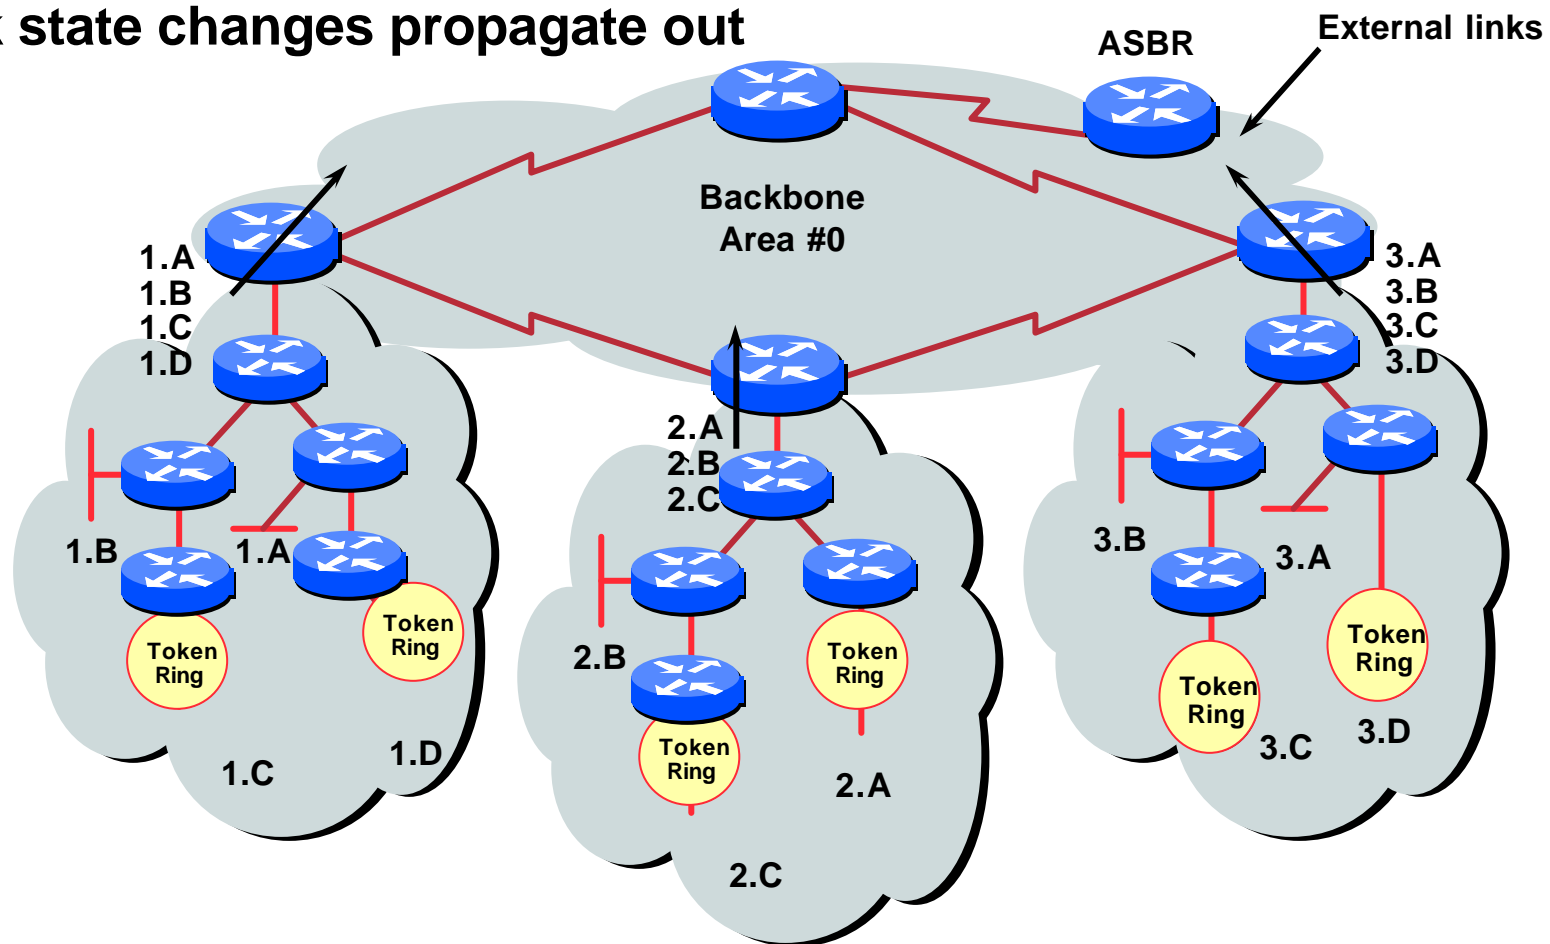
Cisco.com

- **Defines routes to destination external to the AS**
- **Default route is also sent as external**
- **Two types of external LSA:**
 - E1: Consider the total cost up to the external destination**
 - E2: Considers only the cost of the outgoing interface to the external destination**

Not Summarised: Specific Links

Cisco.com

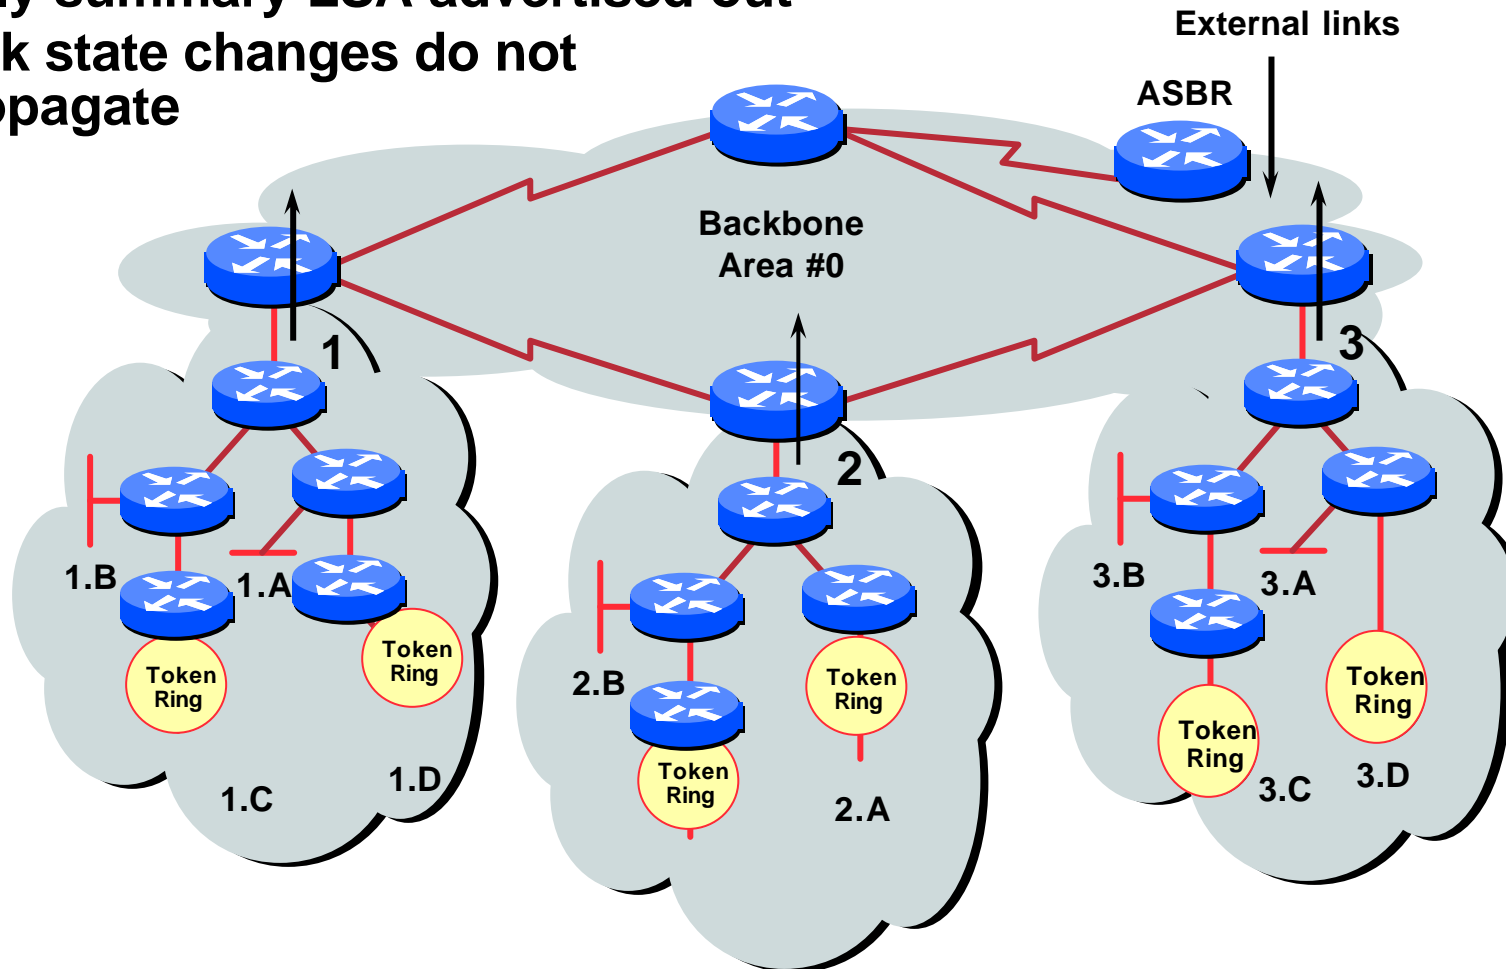
- Specific link LSA advertised out
- Link state changes propagate out



Summarised: Summary Links

Cisco.com

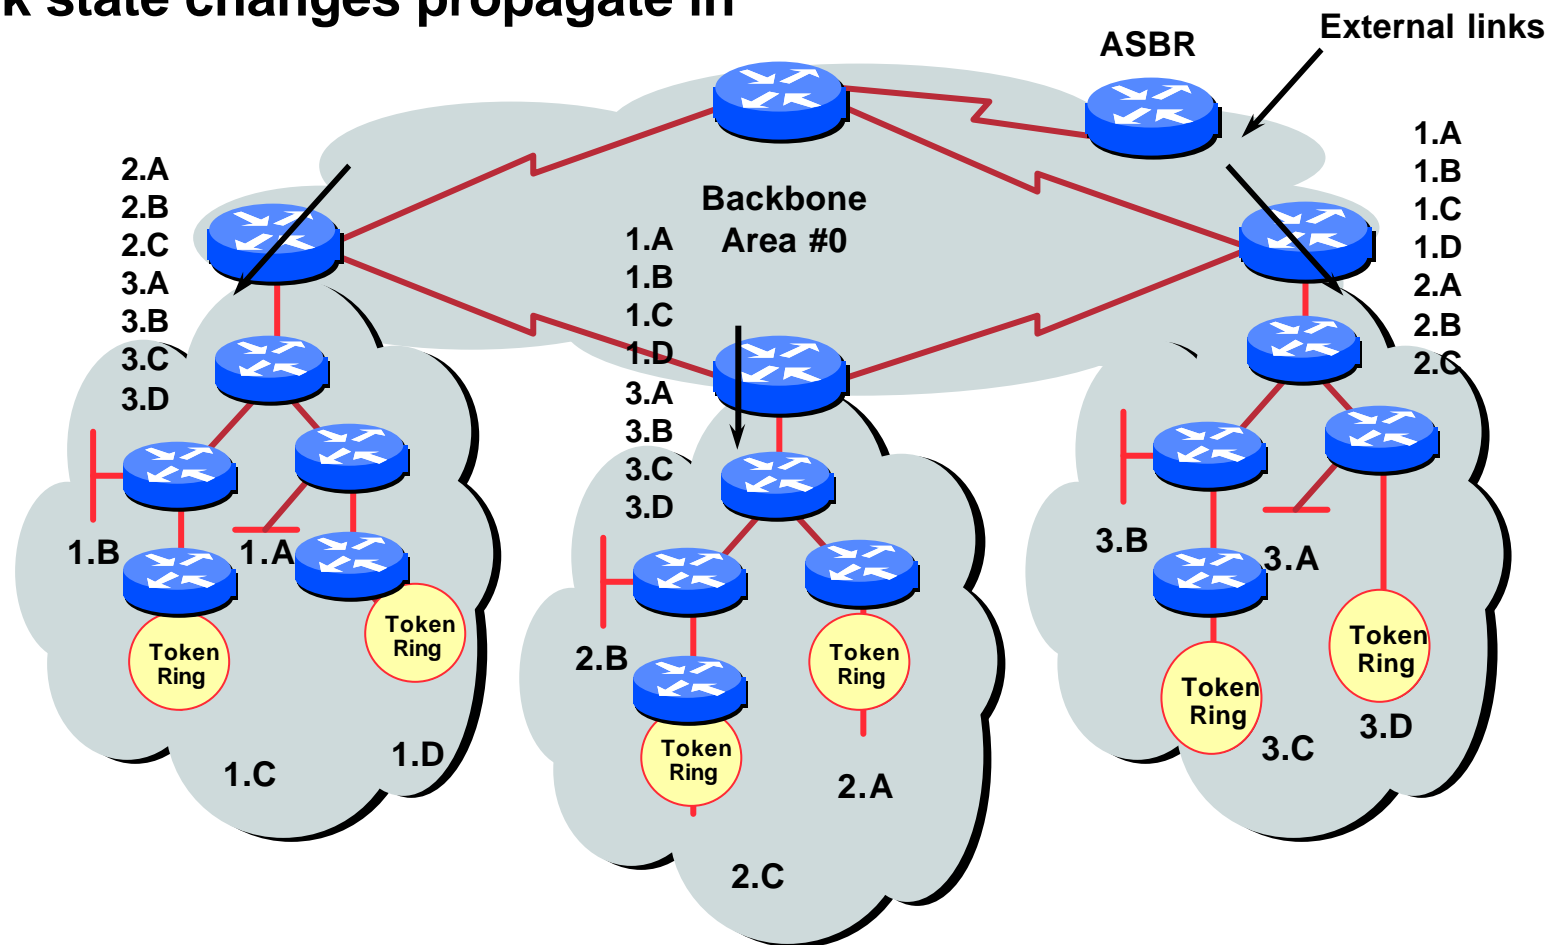
- Only summary LSA advertised out
- Link state changes do not propagate



Not Summarised: Specific Links

Cisco.com

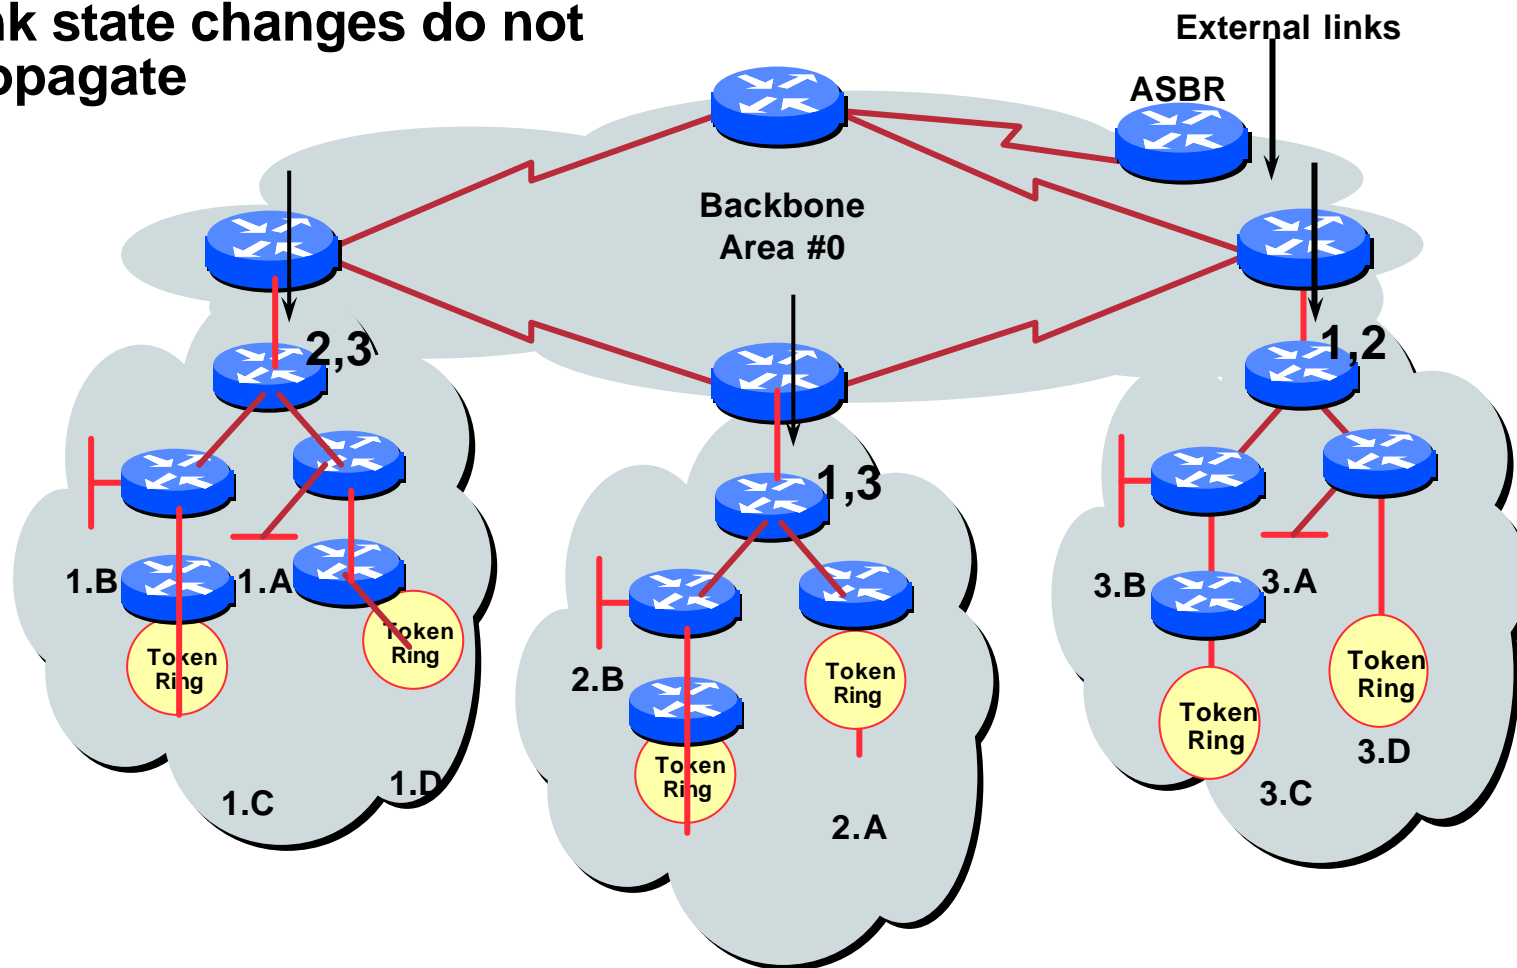
- Specific link LSA advertised in
- Link state changes propagate in



Summarised: Summary Links

Cisco.com

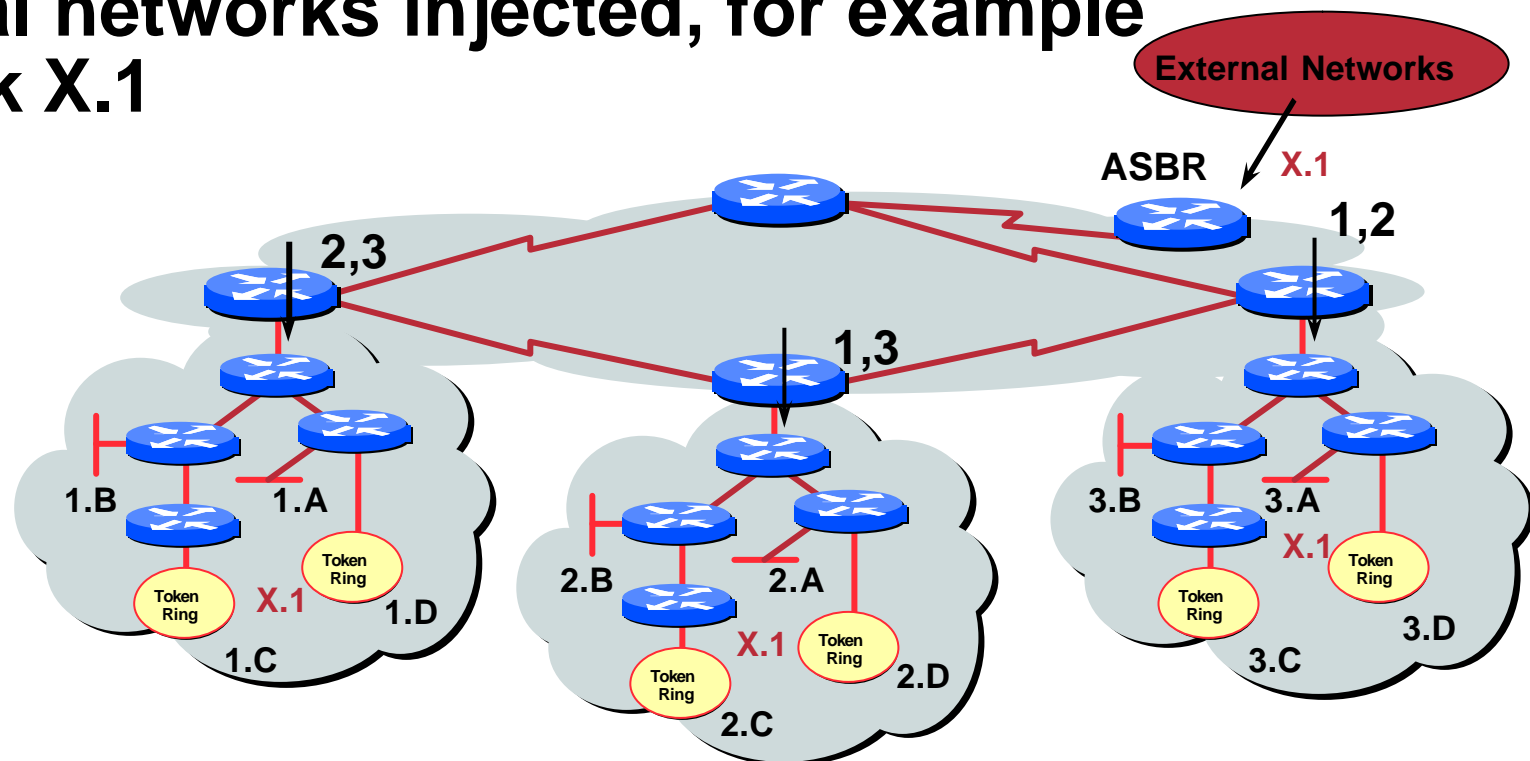
- Only summary LSA advertised in
- Link state changes do not propagate



Regular Area (Not a Stub)

From area 1's viewpoint

- Summary networks from other areas injected
- External networks injected, for example network X.1

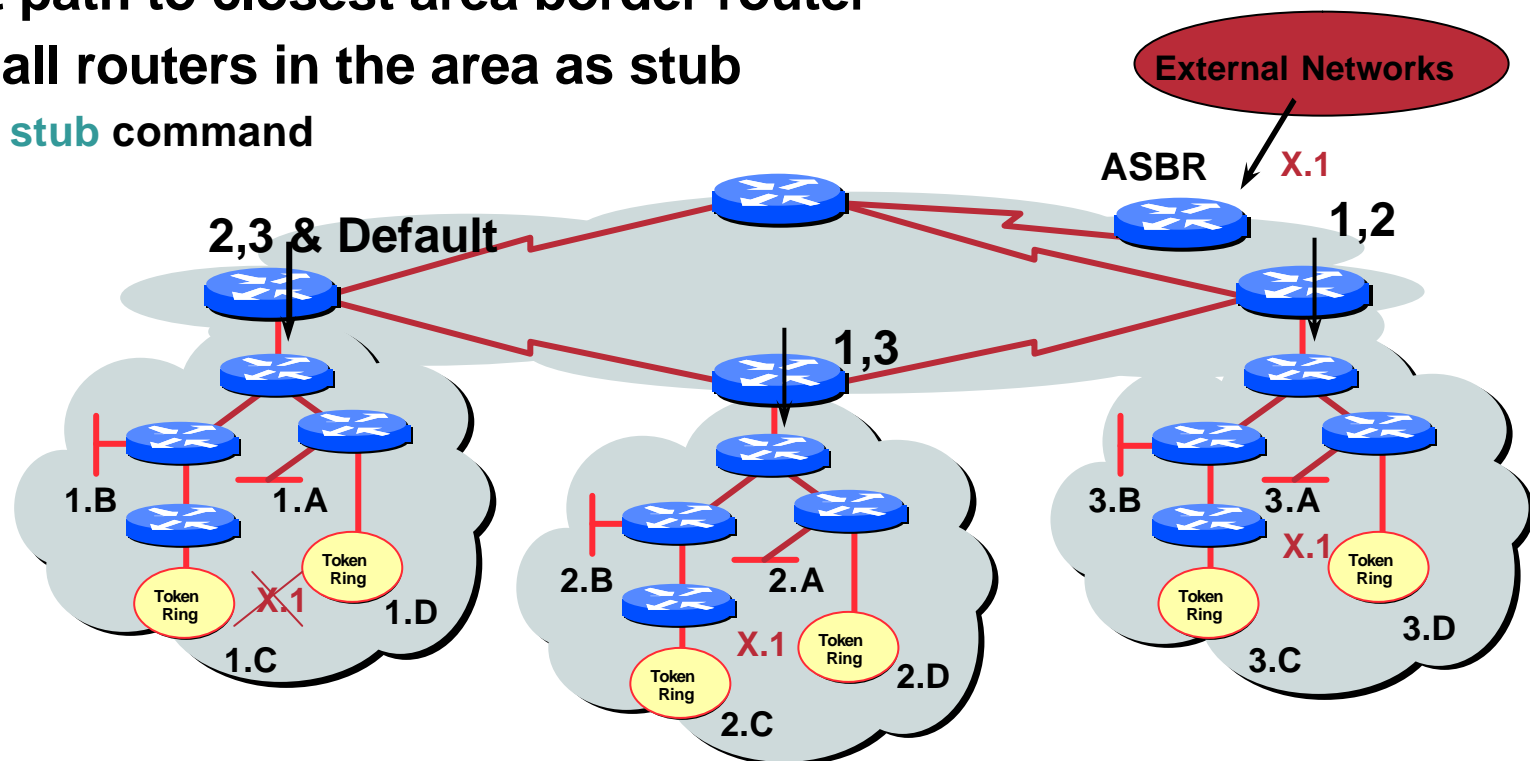


Normal Stub Area

From area 1's viewpoint

- Summary networks from other areas injected
- Default network injected into the area - represents external links
- Default path to closest area border router
- Define all routers in the area as stub

area x stub command

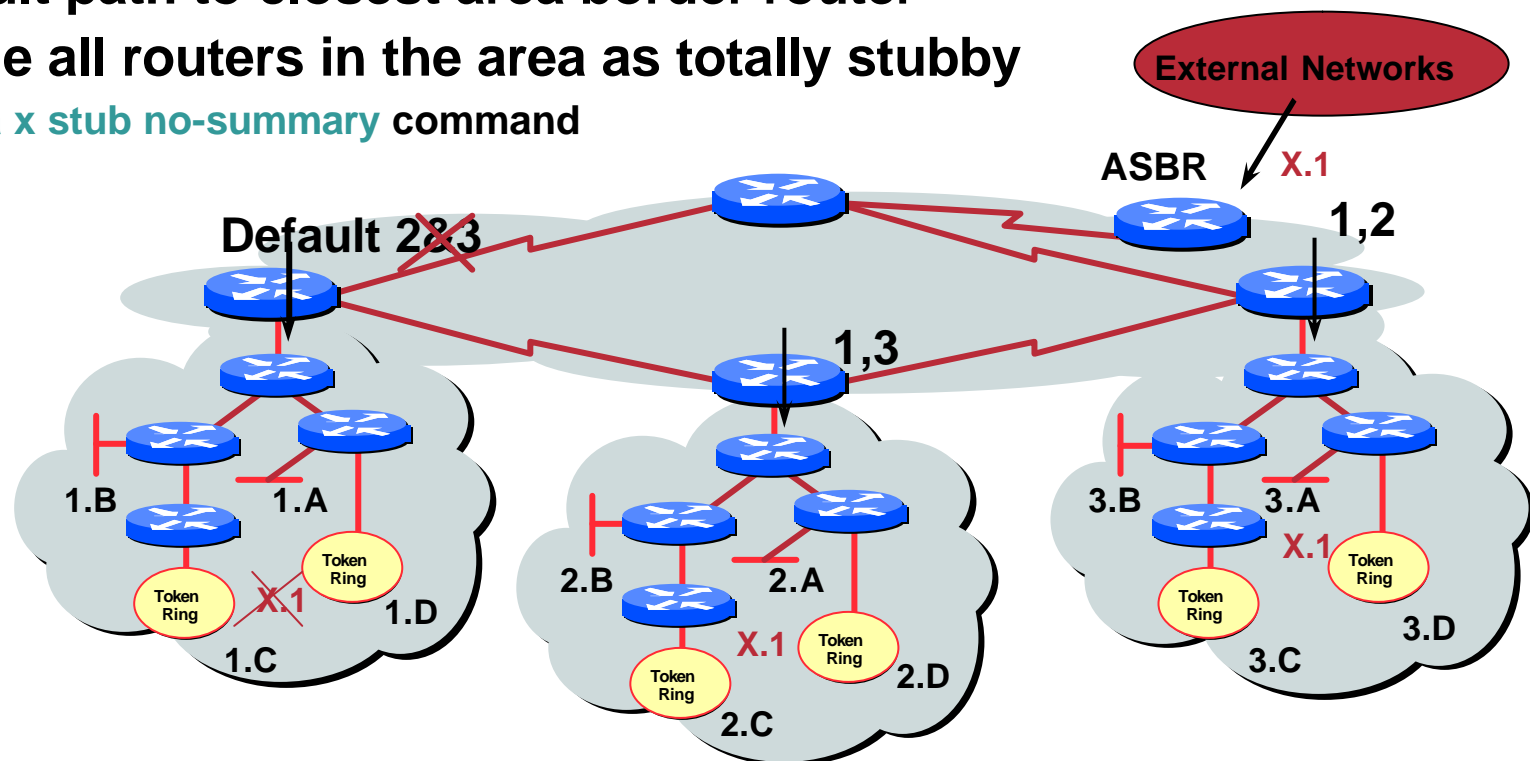


Totally Stubby Area

Cisco.com

From area 1's viewpoint

- Only a default network is injected into the area
Represents external networks and all inter-area routes
- Default path to closest area border router
- Define all routers in the area as totally stubby
`area x stub no-summary` command




Cisco.com

- ### SSA Border routers translate selected type-7 LSAs to type-5 external network LSAs
-
- The diagram illustrates a network topology where SSA border routers translate selected type-7 LSAs to type-5 external network LSAs. The network consists of three internal clouds connected by a central backbone. The left cloud (Cloud 1) has a default router (2,3) and three other routers (1.B, 1.A, 1.C) connected to a Token Ring (1.D). The middle cloud (Cloud 2) has a default router (1,3) and three other routers (2.B, 2.A, 2.C) connected to a Token Ring (2.D). The right cloud (Cloud 3) has a default router (1,2) and three other routers (3.B, 3.A, 3.C) connected to a Token Ring (3.D). External Networks (X.1, X.2) are connected to the default routers. The diagram shows that the default routers (2,3) and (1,2) are translating selected type-7 LSAs to type-5 external network LSAs, as indicated by the red 'X' marks and the text 'Default 2,3' and '1,2'.

Addressing

Cisco.com



Assign contiguous ranges of subnets per area to facilitate summarisation

Summary

Cisco.com

- **Scalable OSPF Network Design**

Area hierarchy

Stub areas

Contiguous addressing

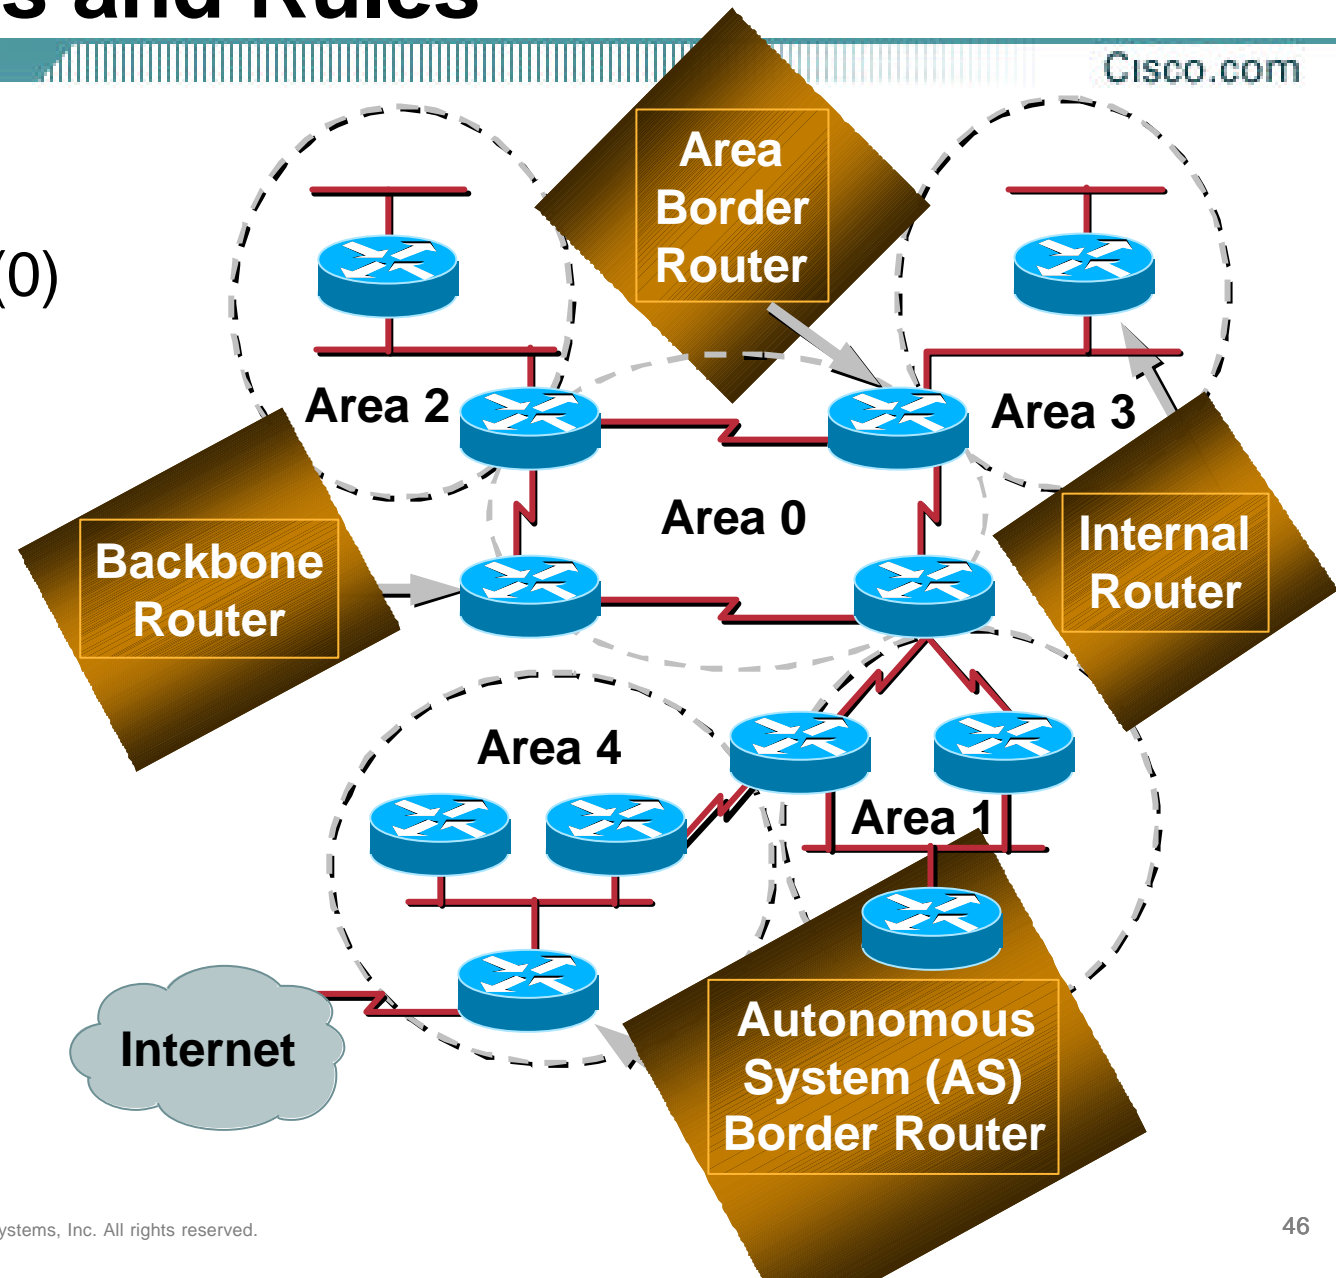
Route summarisation

OSPF Design In Service Provider Networks

OSPF Areas and Rules

Cisco.com

Backbone area (0)
must be present
All other areas
must have
connection
to backbone
Backbone must
be contiguous
Do not partition
area (0)



OSPF Design

Cisco.com

- **OSPF design and Addressing go together.**

**Objective is to keep the Link State Data Base
*lean.***

Create address hierarchy to match topology

**Separate Blocks for infrastructure, customer
interfaces, customers, etc.**

OSPF Design

Cisco.com

- **Examine physical topology**
Is it meshed or hub-and-spoke?
- **Try to use as Stubby an area as possible**
It reduces overhead and LSA counts
- **Push the creation of a backbone**
Reduces mesh and promotes hierarchy

OSPF Design

Cisco.com

- **One SPF per area, flooding done per area**
Watch out for overloading ABRs
- **Different types of areas do different flooding**
 - Normal areas**
 - Stub areas**
 - Totally stubby (stub no-summary)**
 - Not so stubby areas (NSSA)**

OSPF Design

Cisco.com

- **Redundancy**

Dual Links out of each area – using metrics (cost) for traffic engineering

Too much redundancy...

Dual links to backbone in stub areas must be the same – other wise sub-optimal routing will result

Too Much Redundancy in the backbone area without good summarization will effect convergence in the area 0

OSPF for ISPs

Cisco.com

- **OSPF features which should be considered:**

OSPF logging neighbour changes

OSPF reference cost

OSPF Router ID Command

OSPF Process Clear/Restart

OSPF BCP

Adding Networks

OSPF – Adding Networks


Cisco.com

- **BCP – Individual OSPF Network statement for each infrastructure link.**

Have separate IP address blocks for infrastructure and customer networks.

Use *IP Unnumbered* Interfaces or BGP to carry /30s to customers

OSPF should only carry infrastructure routes in an ISP's network.



OSPF – Adding Networks (Method One)

Cisco.com

- **redistribute connected subnets**

Works for all connected interfaces on the router but sends networks as external type-2s – which are not summarized

```
router ospf 100
```

```
    redistribute connected subnets
```

- **Not recommended**

OSPF – Adding Networks

Cisco.com

- **Specific network statements**

Every interface needs a OSPF network statement. Interface that should not be broadcasting OSPF Hello packets needs *passive-interface*.

```
router ospf 100  
  
network 192.168.1.1 0.0.0.3 area 51  
  
network 192.168.1.5 0.0.0.3 area 51  
  
passive interface Serial 1/0
```

OSPF – Adding Networks

Cisco.com

- **Network statements – wildcard mask**

Every interface covered by wildcard mask used in OSPF network statement. Interfaces that should not be broadcasting OSPF Hello packets need *passive-interface* or *default passive-interface*.

```
router ospf 100  
  
network 192.168.1.0 0.0.0.255 area 51  
  
default passive-interface default  
  
no passive interface POS 4/0
```


OSPF – Adding Networks

Cisco.com

- **Key Theme when selecting a technique:
Keep the Link State Database Lean**

Increases Stability

**Reduces the amount of information in the Link
State Advertisements (LSAs)**

Speeds Convergence Time

OSPF – New and Useful Features

OSPF Logging Neighbour Changes

Cisco.com

- The router will generate a log message whenever an OSPF neighbour changes state

- Syntax:

[no] ospf log-adjacency-changes

- Example of a typical log message:

**%OSPF-5-ADJCHG: Process 1, Nbr
223.127.255.223 on Ethernet0 from LOADING to
FULL, Loading Done**

Number of State Changes

Cisco.com

- The number of state transitions is available via **SNMP (ospfNbrEvents)** and the **CLI**:

**show ip ospf neighbor [type number]
[neighbor-id] [detail]**

Detail—(Optional) Displays all neighbours given in detail (list all neighbours). When specified, neighbour state transition counters are displayed per interface or neighbour ID

State Changes (Continued)

Cisco.com

- To reset OSPF-related statistics, use the **clear ip ospf counters EXEC** command. At this point **neighbor** is the only available option; it will reset neighbour state transition counters per interface or neighbour id

clear ip ospf counters [neighbor [<type number>] [neighbor-id]]

OSPF Cost: Reference Bandwidth

Cisco.com

- **Bandwidth used in Metric calculation**

$$\text{Cost} = 10^8 / \text{BW}$$

Not useful for BW > 100 Mbps

- **Syntax:**

```
ospf auto-cost reference-bandwidth <reference-bandwidth>
```

- **Default reference bandwidth still 100 Mbps for backward compatibility**

OSPF Router ID

- If the loopback interface exists and has an IP address, that is used as the router ID in routing protocols – **stability!**
- If the loopback interface does not exist, or has no IP address, the router ID is the highest IP address configured – **danger!**
- New sub command to manually set the OSPF Router ID:

```
router-id <ip address>
```

OSPF Clear/Restart

- **clear ip ospf [pid] redistribution**

This command can now clear redistribution based on OSPF routing process ID. If no pid is given, it assumes all OSPF processes.

- **clear ip ospf [pid] counters**

This command can now clear counters based on OSPF routing process ID. If no pid is given, it assumes all OSPF processes.

- **clear ip ospf [pid] process**

This command will restart the specified OSPF process. If no pid is given, it assumes all OSPF processes. It attempts to keep the old router-id, except in cases, where a new router-id was configured, or an old user configured router-id was removed. Since this command can potentially cause a network churn, a user confirmation is required before performing any action.

OSPF Command Summary

Redistributing Routes into OSPF

Cisco.com

```
ROUTER OSPF <pid#x>  
  REDISTRIBUTE {protocol} <as#y>  
    <metric>  
    <metric-type (1 or 2)>  
    <tag>  
    <subnets>
```

Router Sub-commands

Cisco.com

- **NETWORK <n.n.n.n> <mask> AREA <area-id>**
- **AREA <area-id> STUB {no-summary}**
- **AREA <area-id> AUTHENTICATION**
- **AREA <area-id> DEFAULT_COST <cost>**
- **AREA <area-id> VIRTUAL-LINK <router-id>...**
- **AREA <area-id> RANGE <address mask>**

Interface Subcommands

Cisco.com

- **IP OSPF COST <cost>**
- **IP OSPF PRIORITY <8-bit-number>**
- **IP OSPF HELLO-INTERVAL <number-of-seconds>**
- **IP OSPF DEAD-INTERVAL <number-of-seconds>**
- **IP OSPF AUTHENTICATION-KEY <8-bytes-of-password>**

Introduction to OSPF

ISP/IXP Workshops