NETWORKERS 2003

THE POWER TO TRANSFORM BUSINESS. now.

Cisco Systems

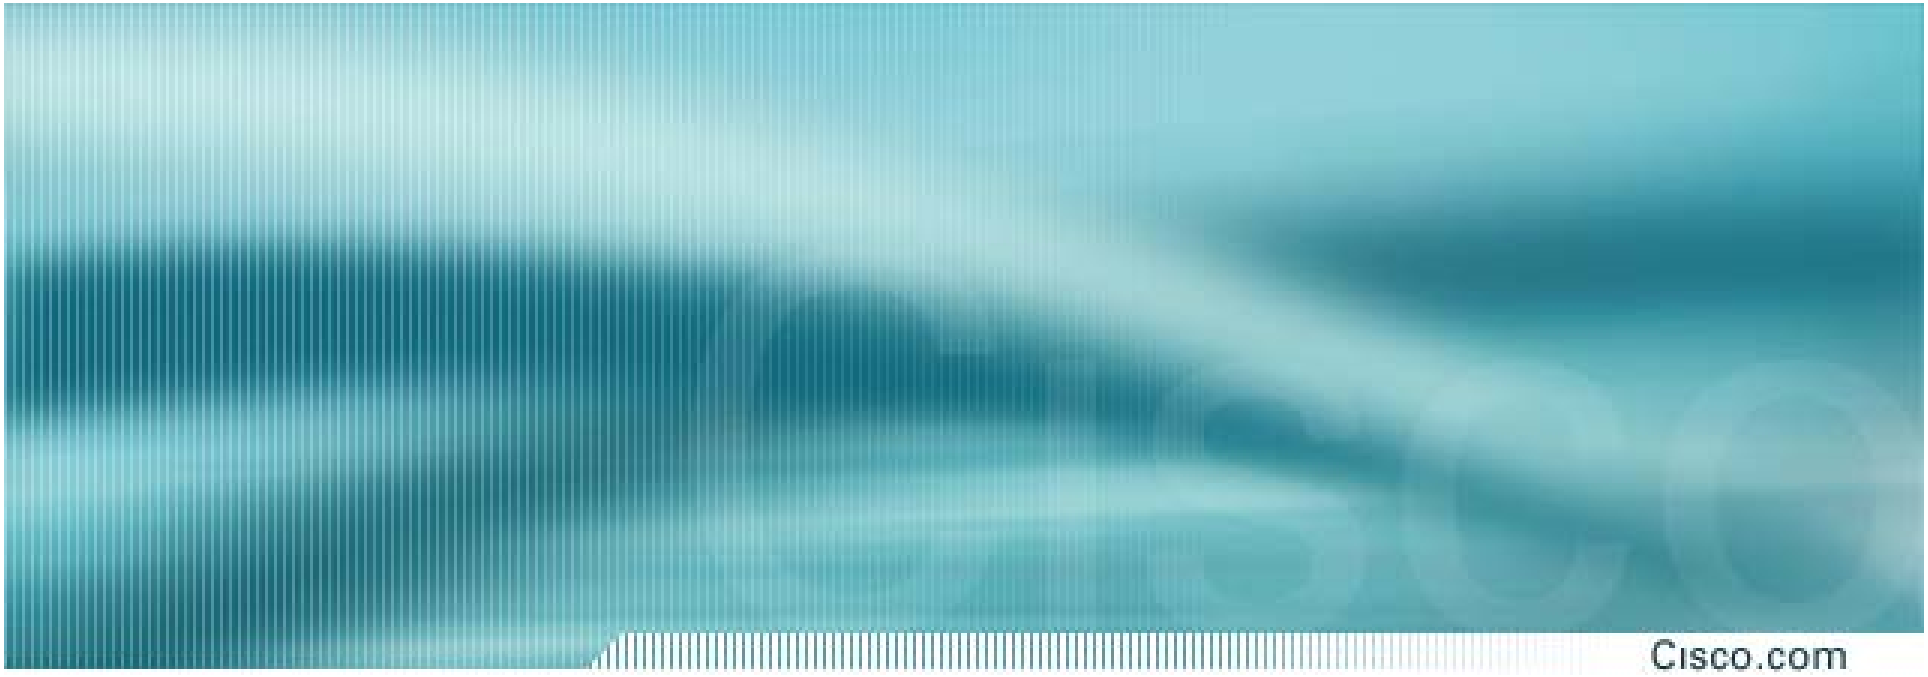

Deployment and Analysis of Link State Protocols

Session RST-2002

Agenda

Cisco.com

- **Pre-Requisite**
- **Link State Protocol Refresher**
- **Fundamental Deployment and Analysis**
- **OSPF Configuration Examples**
- **ISIS Configuration Examples**

Pre-Requisite

- **OSPF and ISIS basic concepts**
- **Some experience with OSPF and ISIS**
- **RST-1001**

Agenda

- Pre-Requisite
- Link State Protocol Refresher
- Fundamental Deployment and Analysis
- OSPF Configuration Examples
- ISIS Configuration Examples

Link State Protocols

- **Two popular link state IGPs**

  OSPF and ISIS

- **Hierarchical model**
- **Same topology within an area**
- **SPF algorithm**

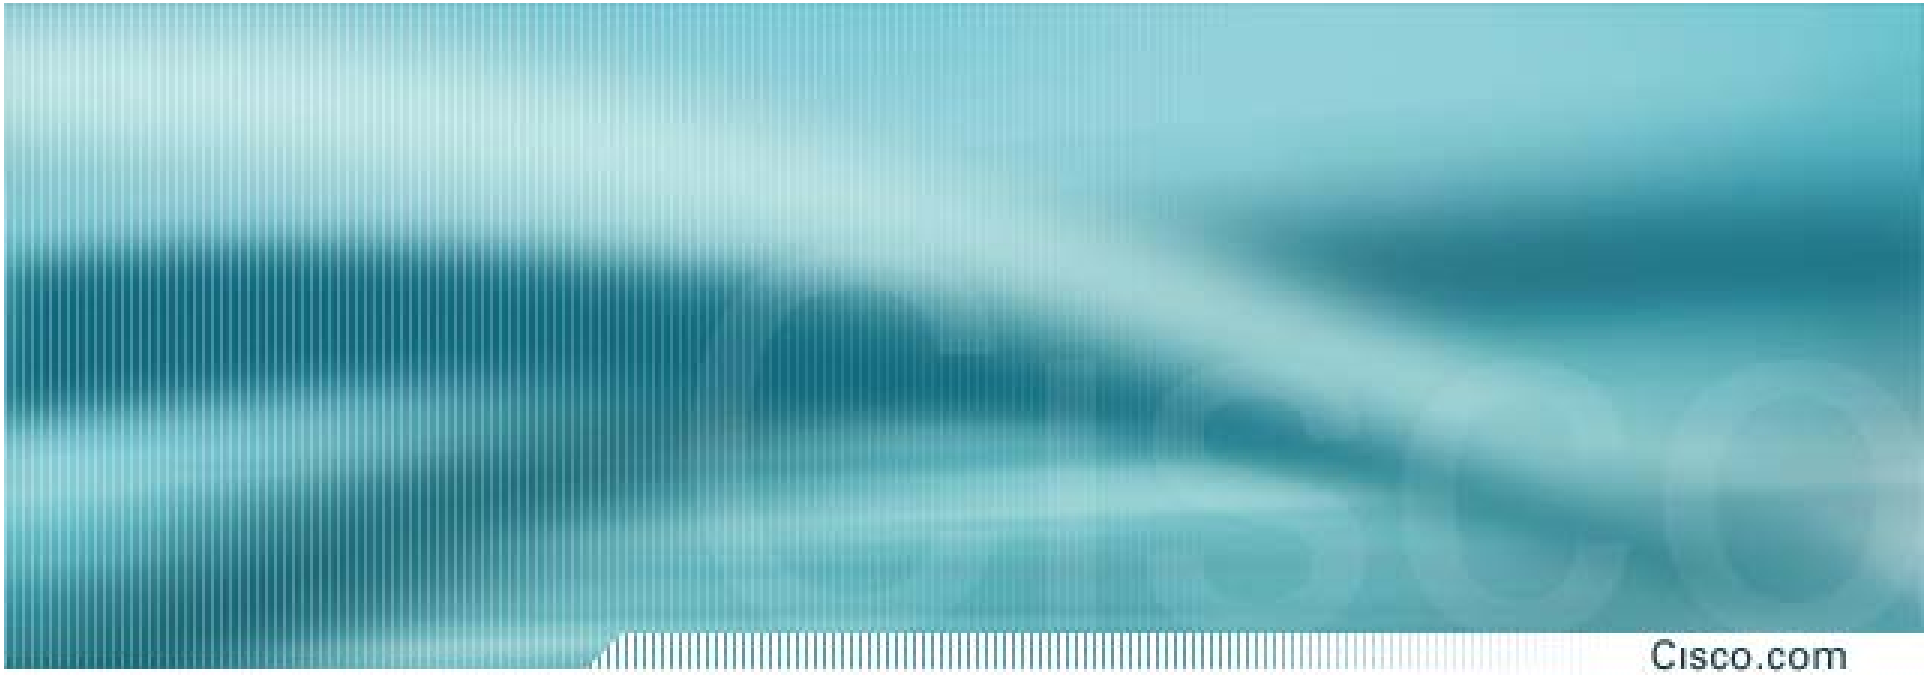

IS-IS Topology Basics

Level-1 Routing

- L1-only routers know only topology of their own area (including all ISs and ESs in the area)
- Traffic to other areas is sent via the closest L2 IS
- L1L2 ISs set the "attached-bit" in their L1 LSP

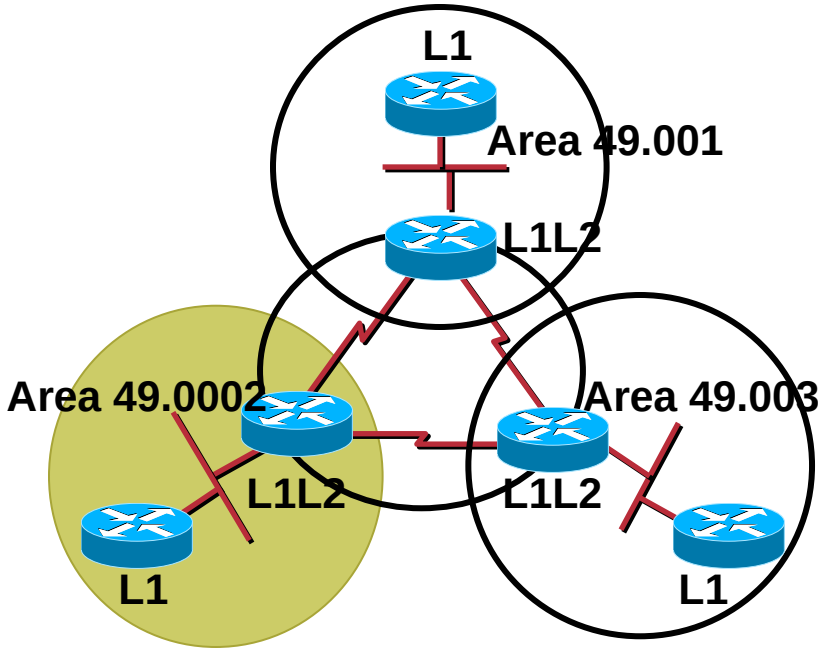

Level-2 Routing

- **L2 routers know about other areas**
  - **L2 area addresses and L2 routers**
- **When doing OSI routing, the L2 ISs must know their own area; therefore never use L2-only on OSI routers**
  - **L2-only is possible when doing just IP**

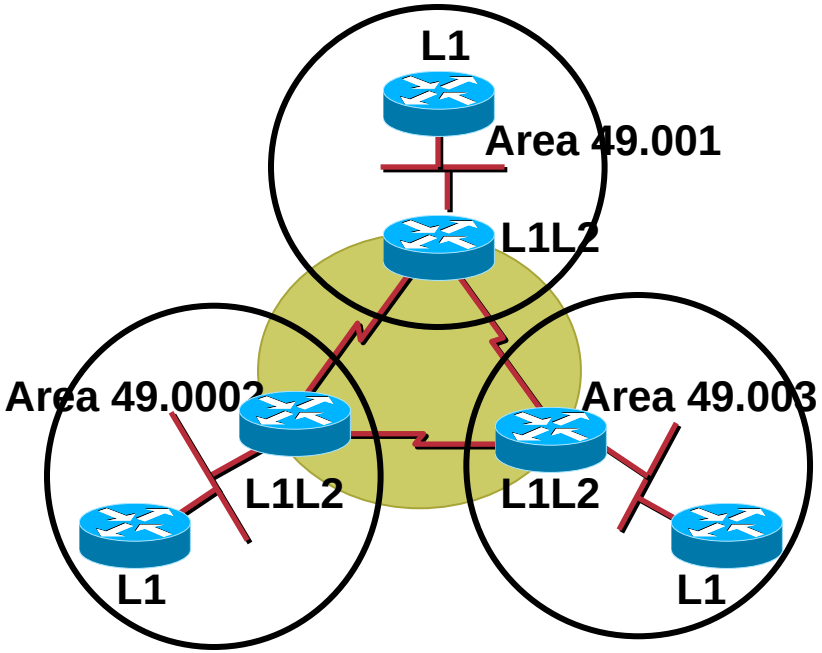

Level-2 Routing

- **Transit traffic requires routers inside the area to know about other areas**
  - **Routers in transit paths must be L1L2 routers to have the full L2 LSDB**
  - **Similar to pervasive BGP requirement**
- **L2 routers must be contiguous**

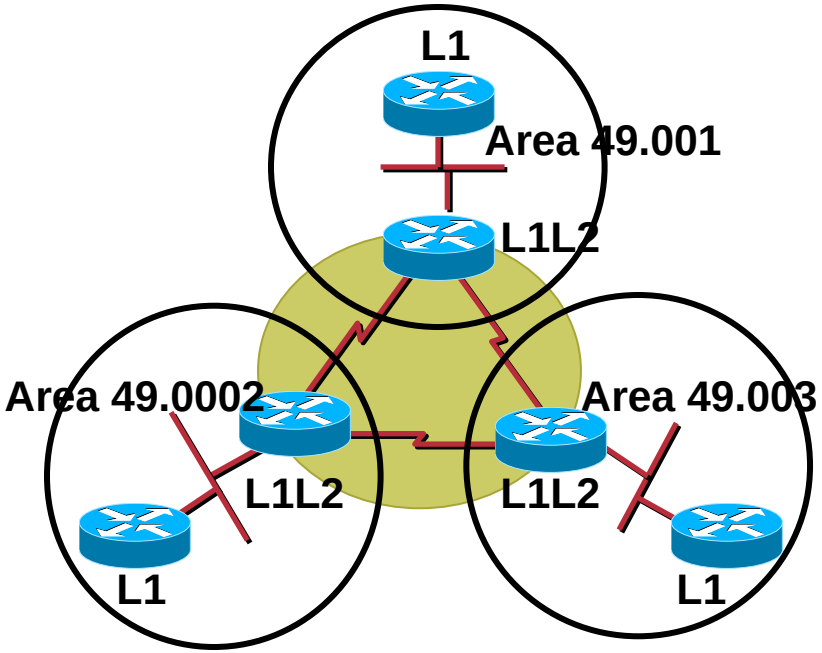

The Backbone

- **A router can't tell whether it is a transit IS without help of a human**

  Therefore the Cisco default is to be L1L2; this will make the backbone larger then necessary so always configure L1-only when possible

- **L1L2 in one area is less scalable**

  Especially with ISIS for IP

Areas and Backbone Routers

Cisco.com

- **Backbone must be L2 contiguous**

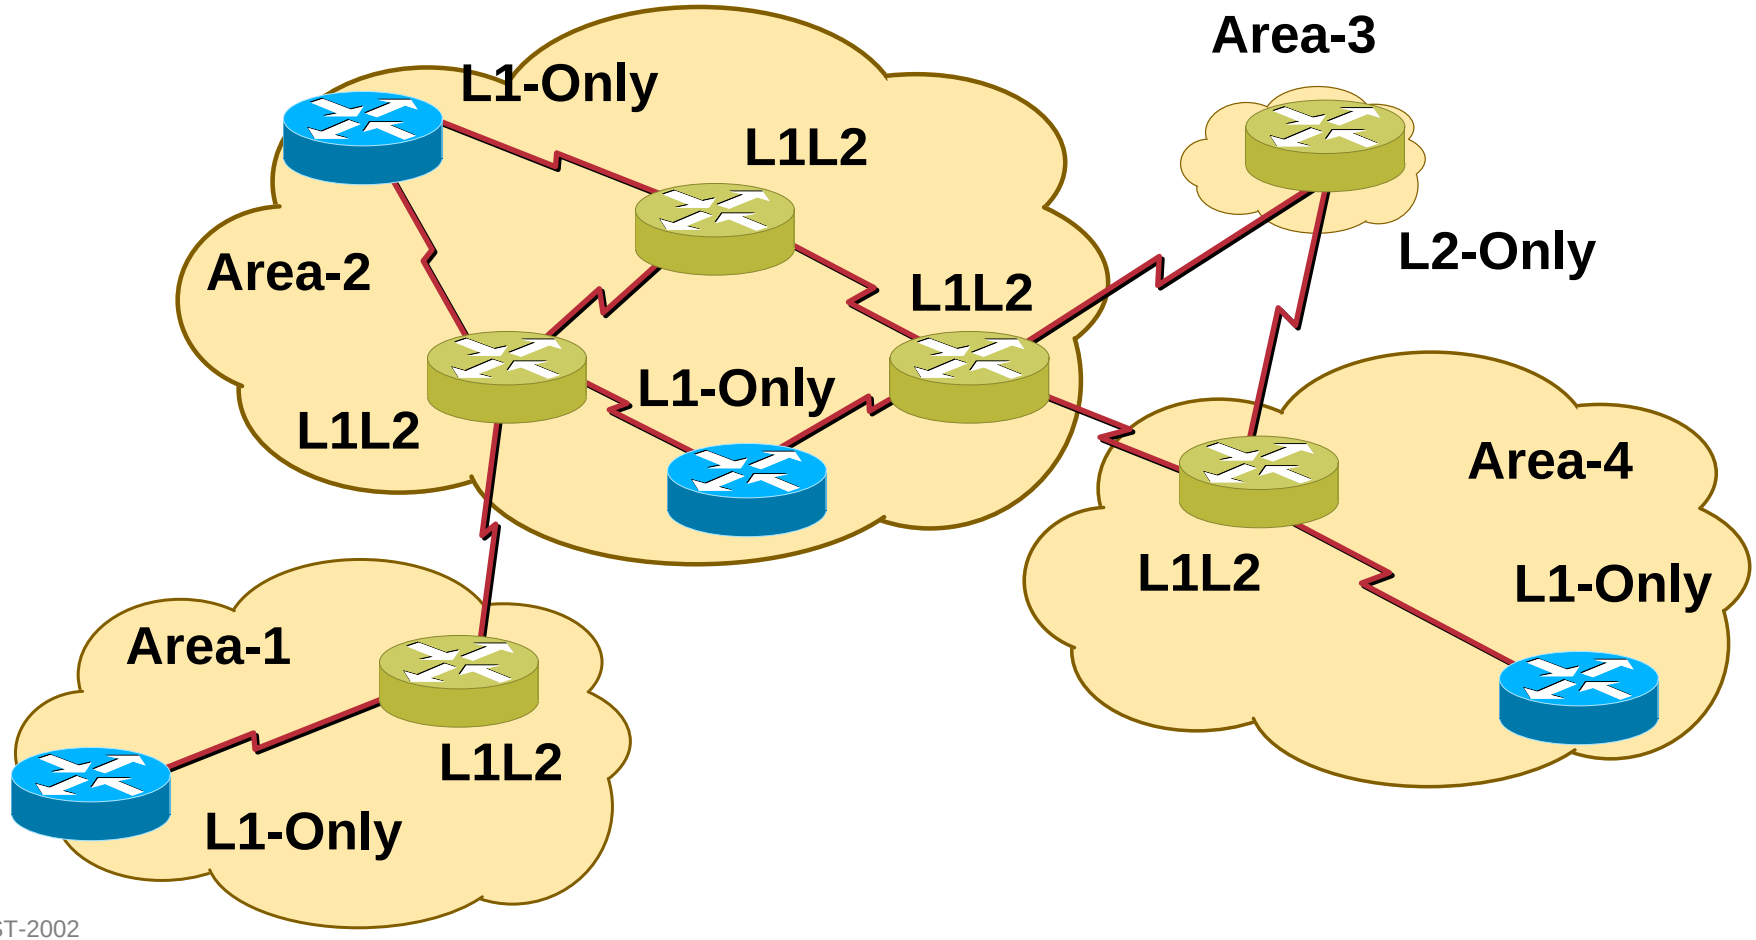

Areas and Backbone Routers

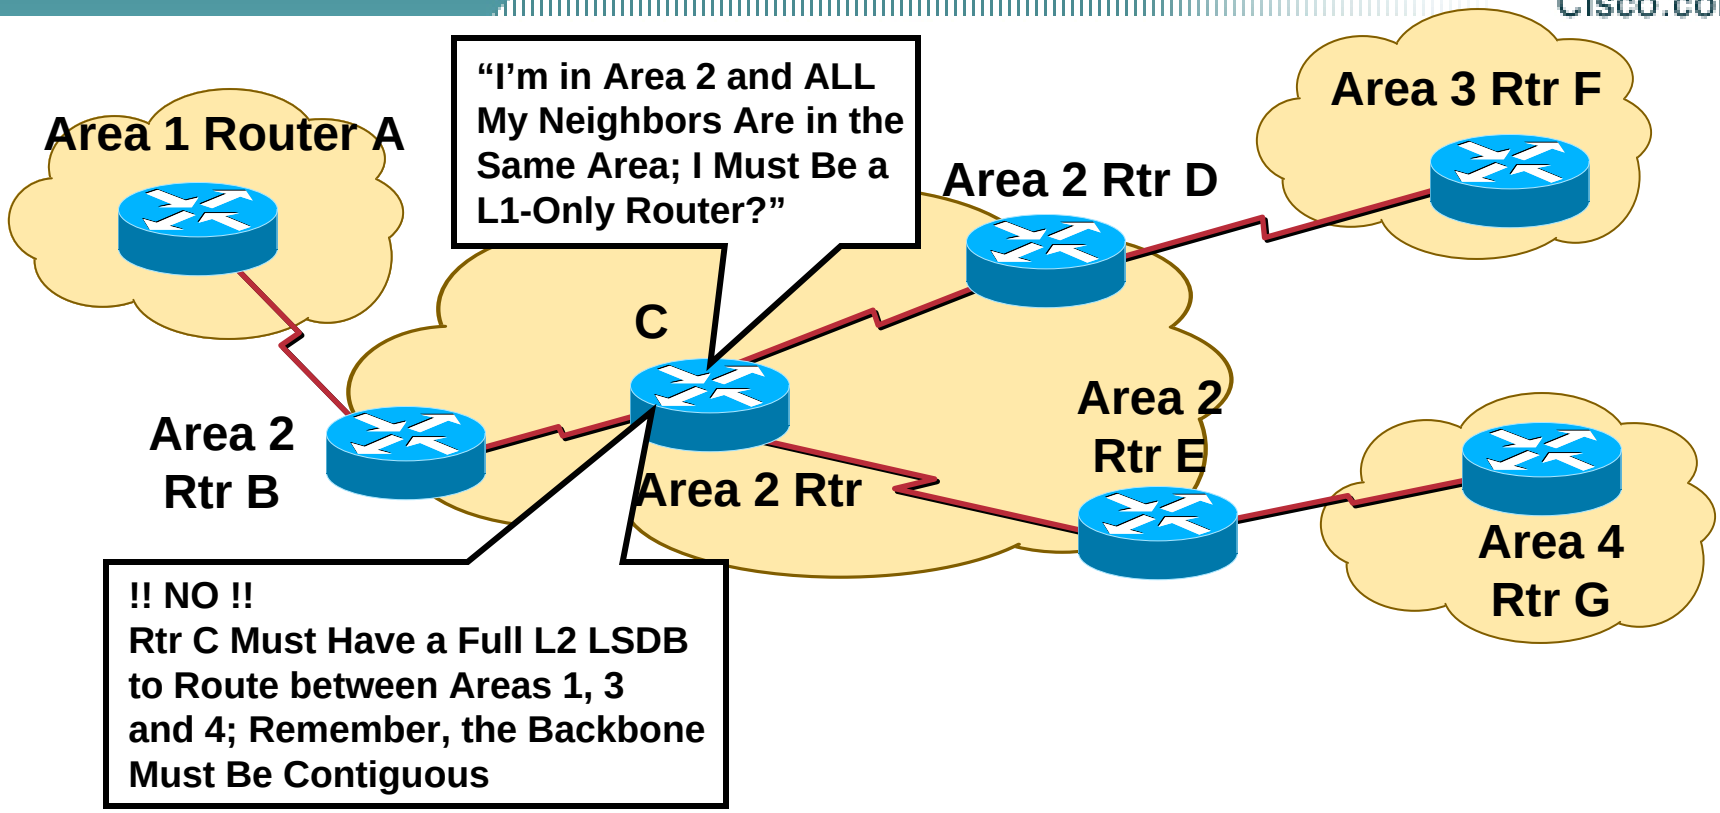

Remember, the Backbone Must Be Contiguous

ISIS Router Cannot Determine if They Need To Be L1 or L1L2
So All Routers Try To Be an L1L2 IS by Default

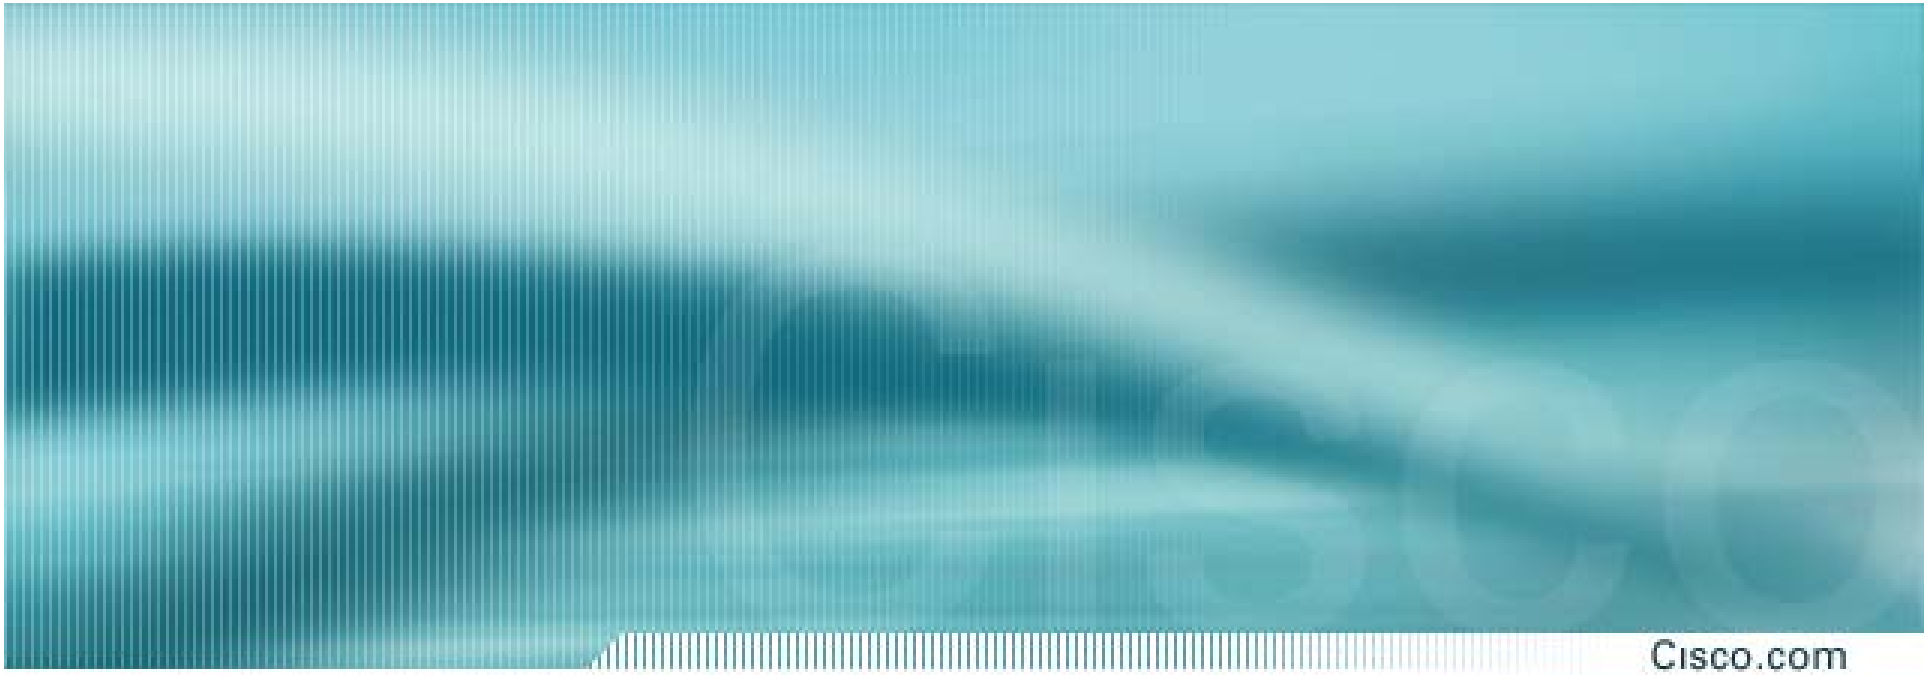

OSPF Topology Basics

OSPF Areas

- OSPF uses a 2 level hierarchical model
- Areas defined with 32 bit number
  - Defined in IP address format
  - Can also be defined using single decimal value (i.e., Area 0.0.0.0, or Area 0)
- 0.0.0.0 reserved for the backbone area

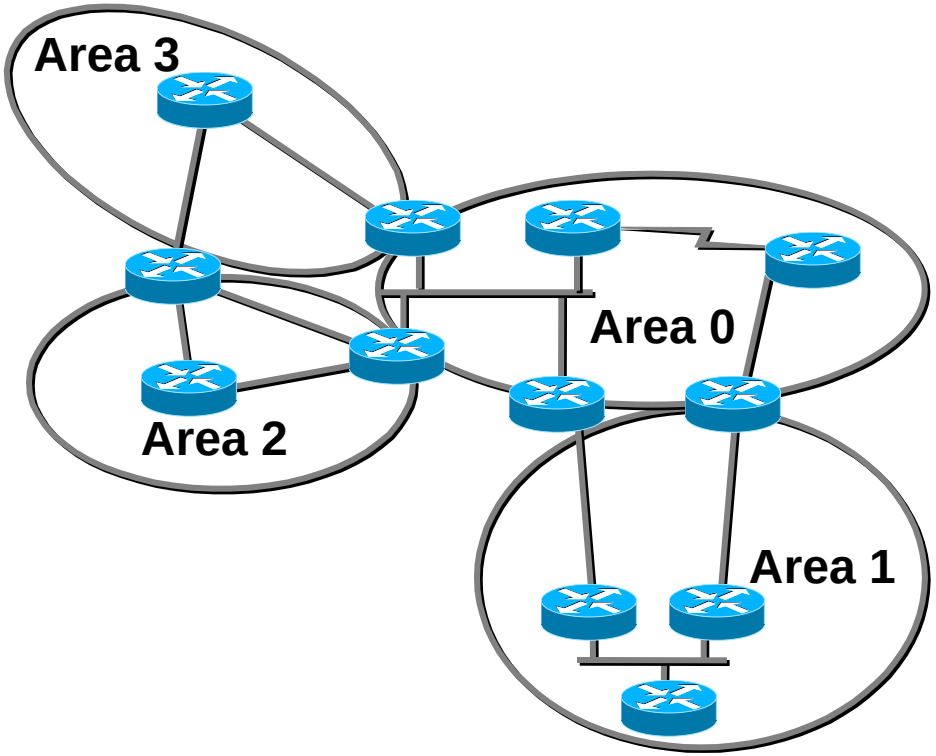

OSPF LSAs

- Router and network LSA's within an area
- Summary LSA type 3 outside the area
- Summary LSA Type 4 and Type 5 for redistributed routes
- Partial SPF for summary and External LSA's

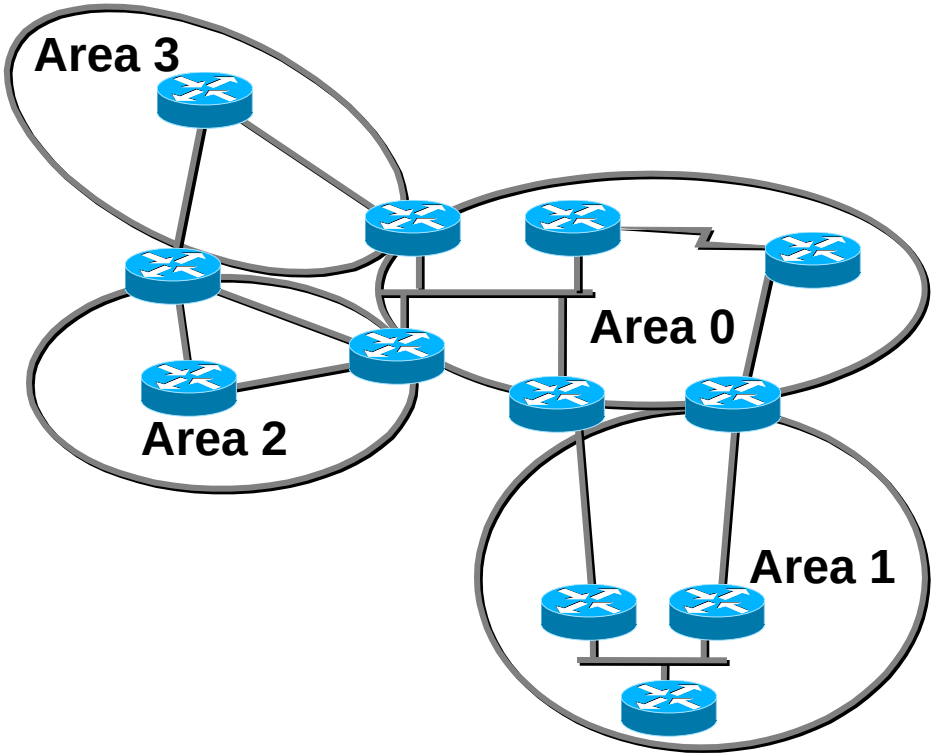

OSPF Virtual Links

- **Can be useful for several purposes**
  - **Allow areas to connect to areas other than 0**
  - **Repair a discontinuous area 0**
  - **Optimal path purpose**
  - **Backup purpose**

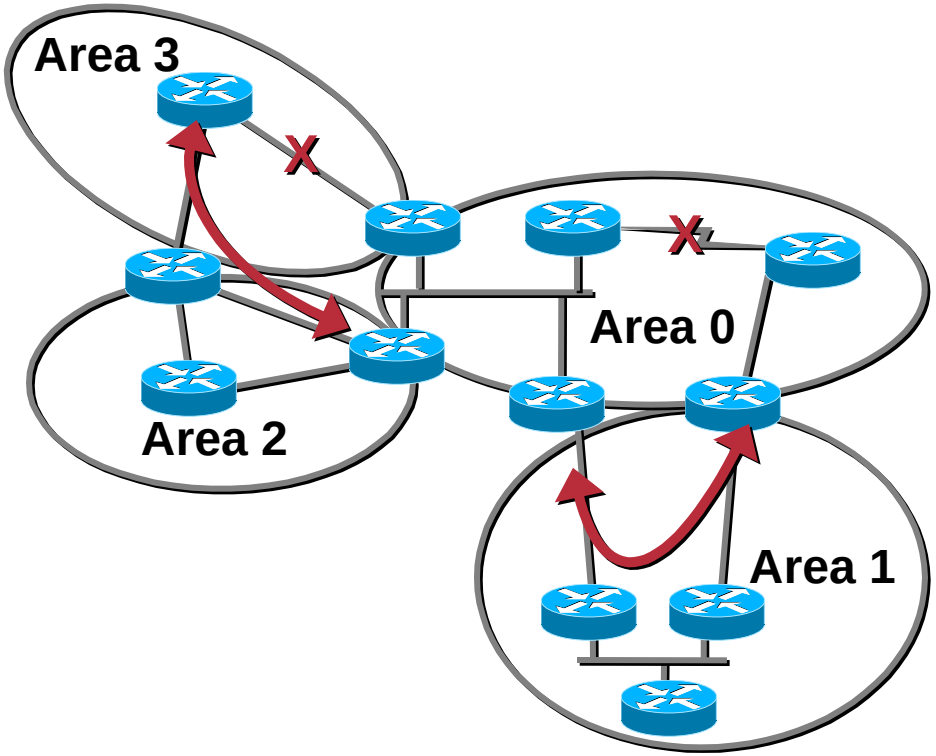

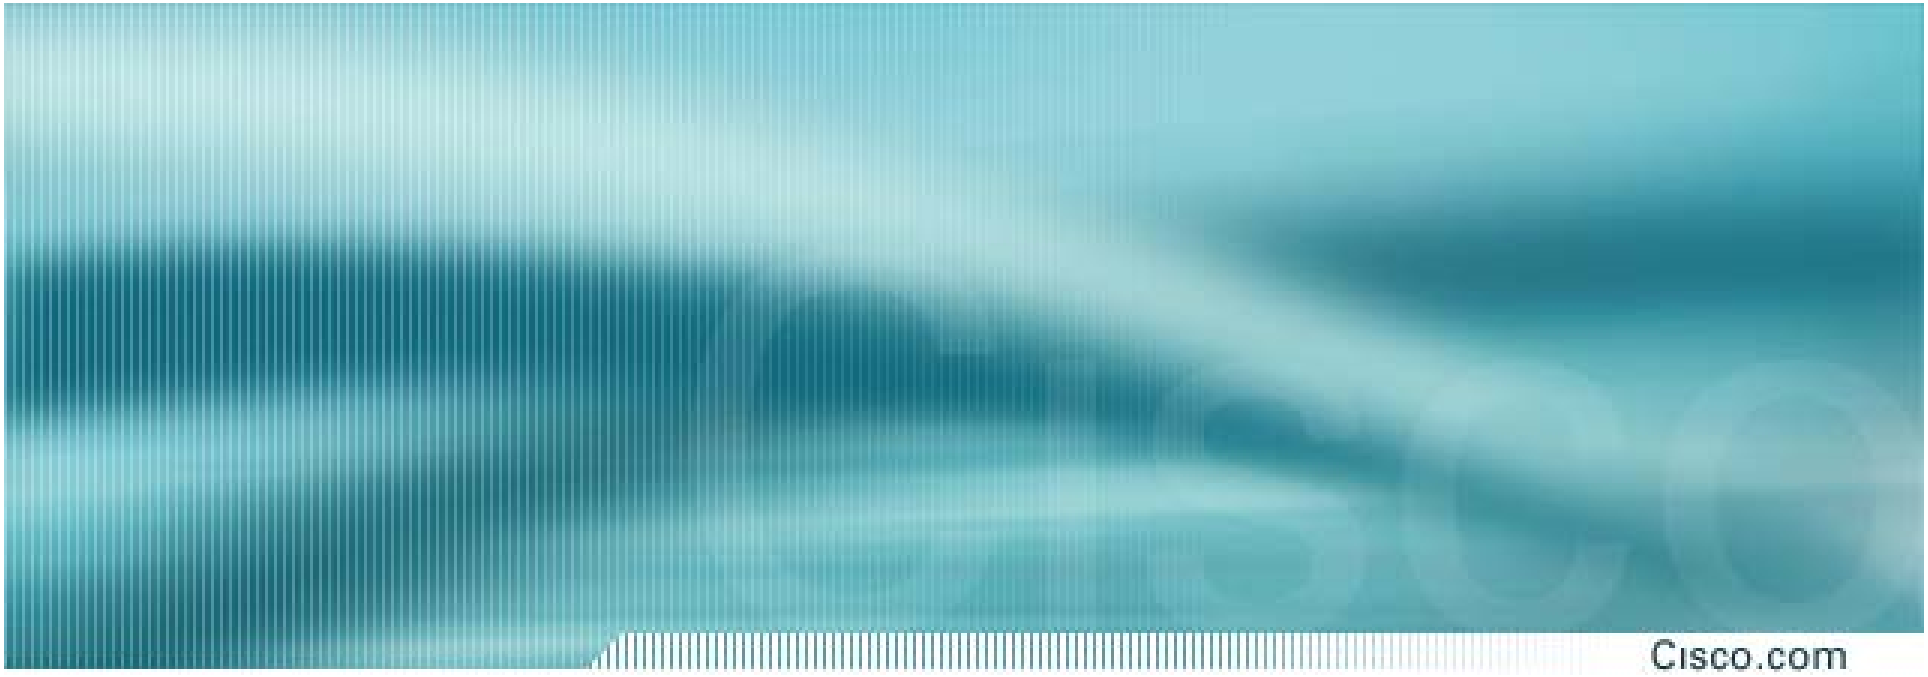

Scaling OSPF

Summarization

- Routing is by longest prefix match
- Instead of advertising many more specific prefixes, advertise only one summary prefix
  - Area-range on ABR to summarize type 3 LSAs
  - Summary-address on ASBR to sum type 5s
- Not only smaller, but also more stable
- Drawback is possible suboptimal routing

Not Summarized: Specific Links

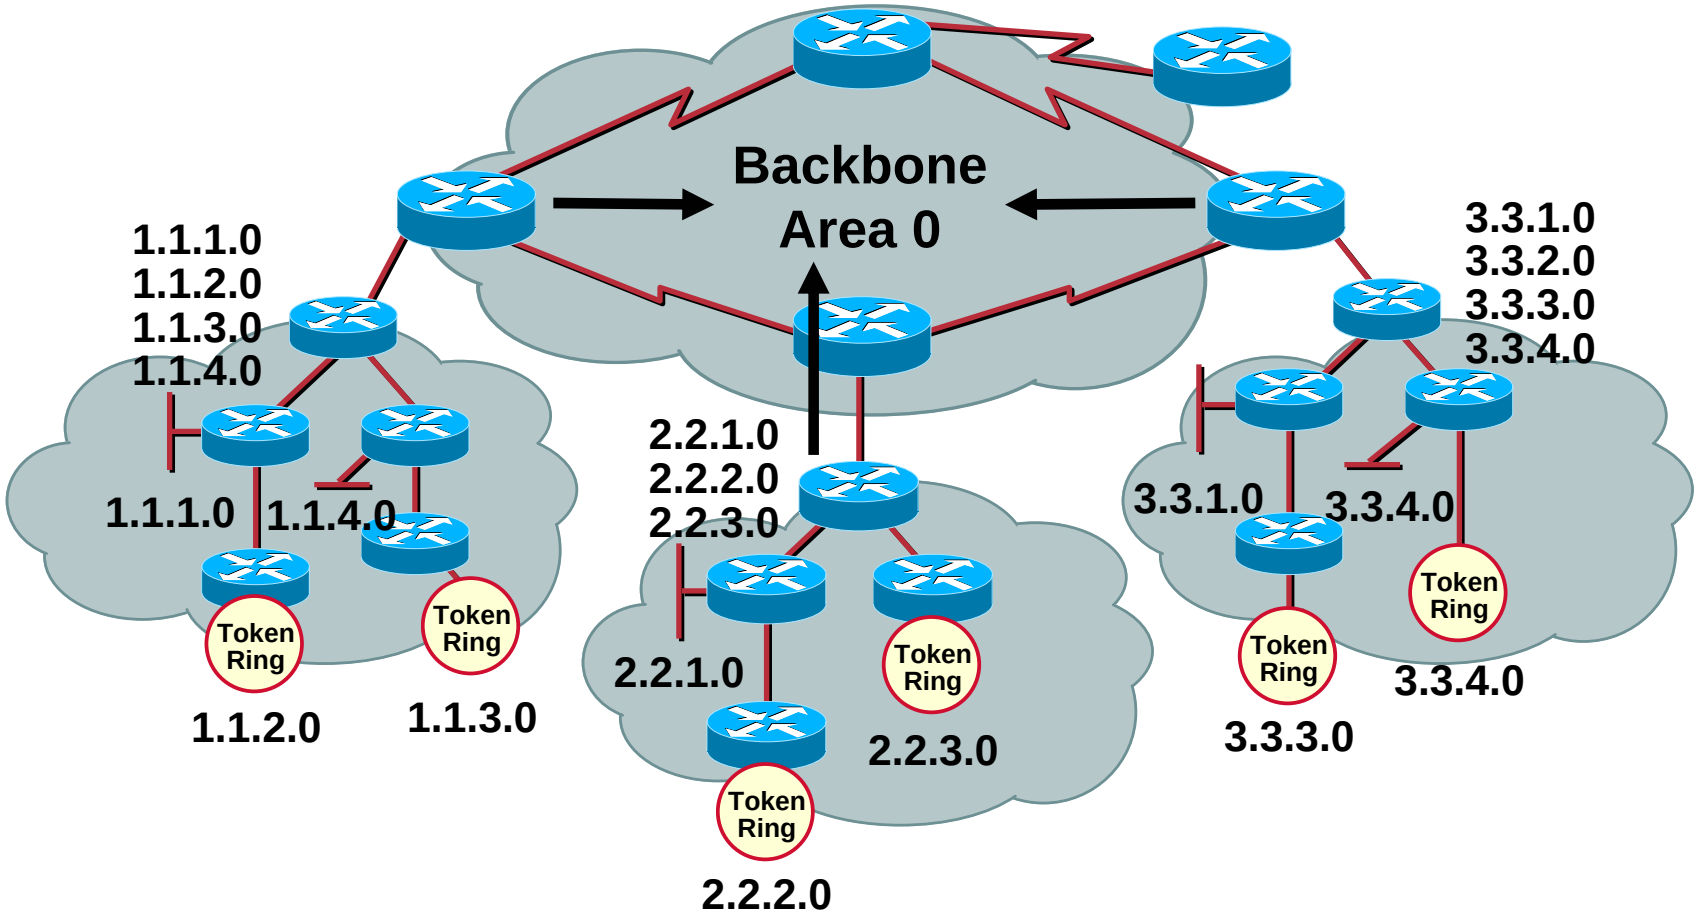

- **Only summary LSA advertised out**
- **Link-state changes do not propagate**

Summarized: Summary Links

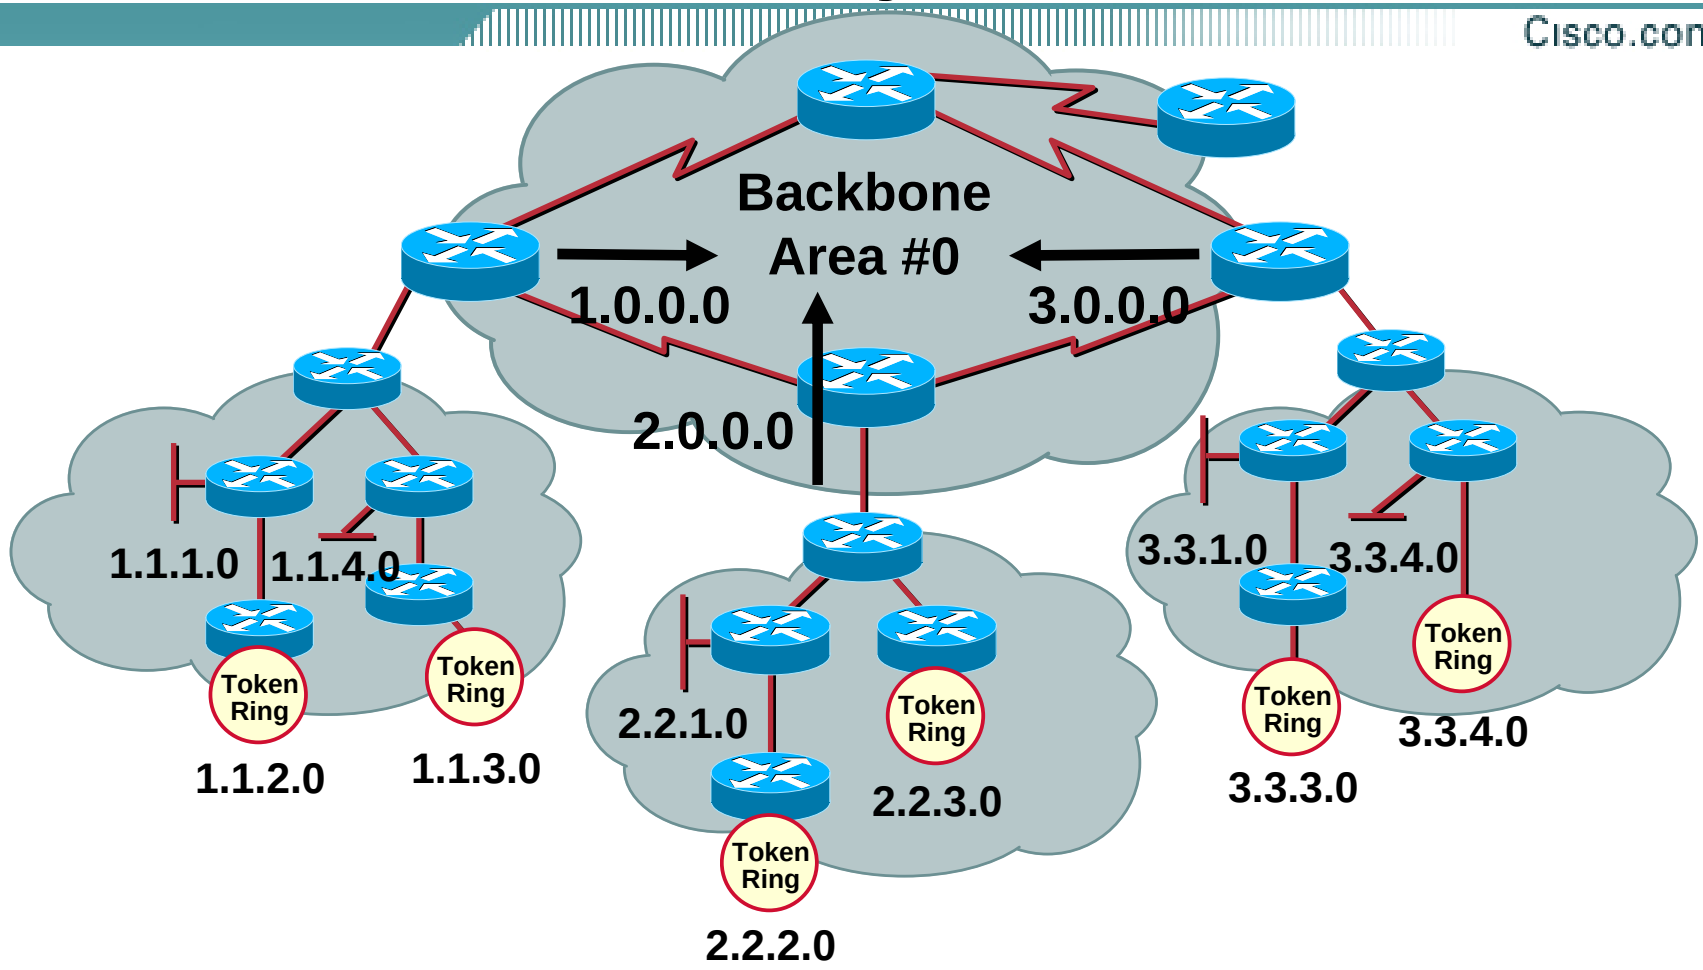

- Only summary LSA advertised out
- Link-state changes do not propagate

Summarization (Cont.)

Configure on Both ABRs
Area-Range 11.1.0/17
Area-Range 11.1.128/17

Cost to Range 1:
Via ABR1: 30
Via ABR2: 80

Cost to Range 2:
Via ABR1: 80
Via ABR2: 30

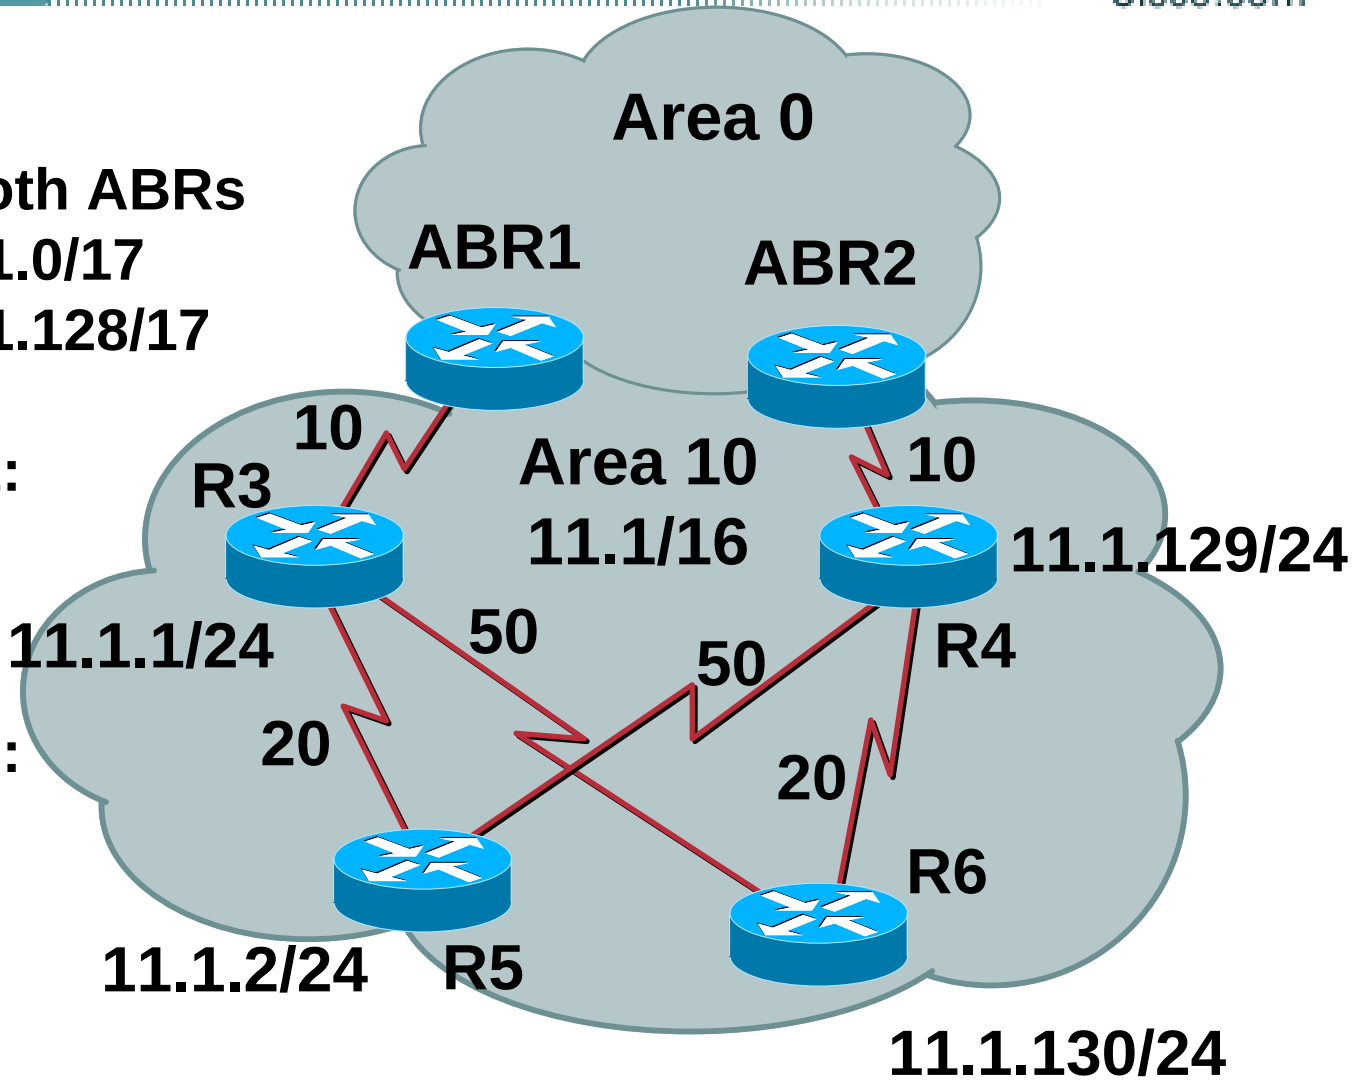

Using Areas

- The tool to make OSPF scale
- One SPF per area, flooding done per area
- Different types of areas do different flooding
  - Normal areas
  - Stub areas
  - Totally stubby (stub no summary)
  - Not so stubby areas (NSSA)

Regular Area

- **Regular areas**
  - **Summary LSA (summarized/non-summarized) from other areas injected**
  - **External links injected**

Regular Area

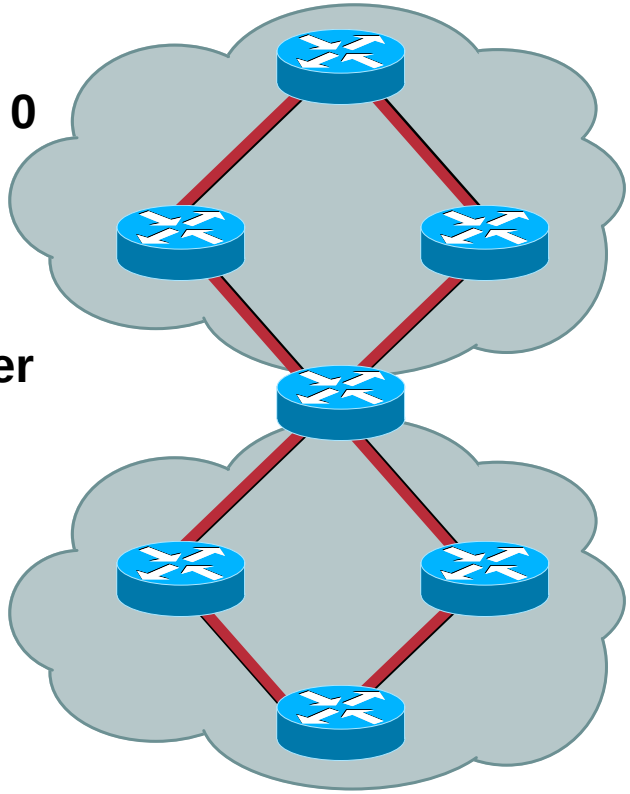

An Area Receives the Following

- Summary LSAs
- Specific LSAs
- Default External 0.0.0.0
- Summarized/Non-summarized Specific External Links

- **ABRs forward all LSAs from backbone**

Stub Area

- **Stub area**
  - **Summary LSAs from other areas injected**
  - **LSA type 5 not injected**
  - **Default LSA injected into area as summary LSA**
  - **Define all routers in area as stub**
  - **External link flaps will not be injected**

Stub Area

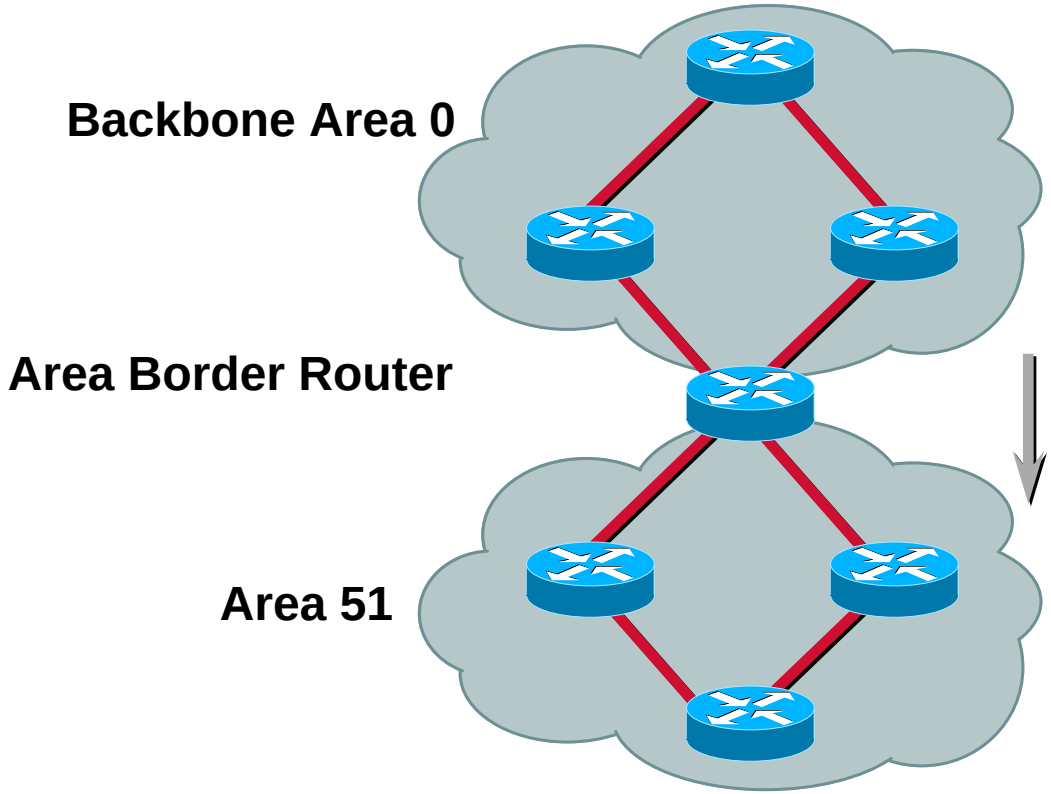

Type 5 LSA Filtering

A Stub Area Receives

- **Summary LSAs**
- **Specific LSAs**
- **Default External 0.0.0.0**

- **Consolidates specific external links—default 0.0.0.0**

Totally Stubby Area

- **Totally stubby area**
  - **Default LSA injected into area**
  - **Represents all external links**
  - **Represents all summarized internal links**
  - **Represents non-summarized internal links**
  - **Very stable, small LSDB, fewer routes**

Totally Stubby Area

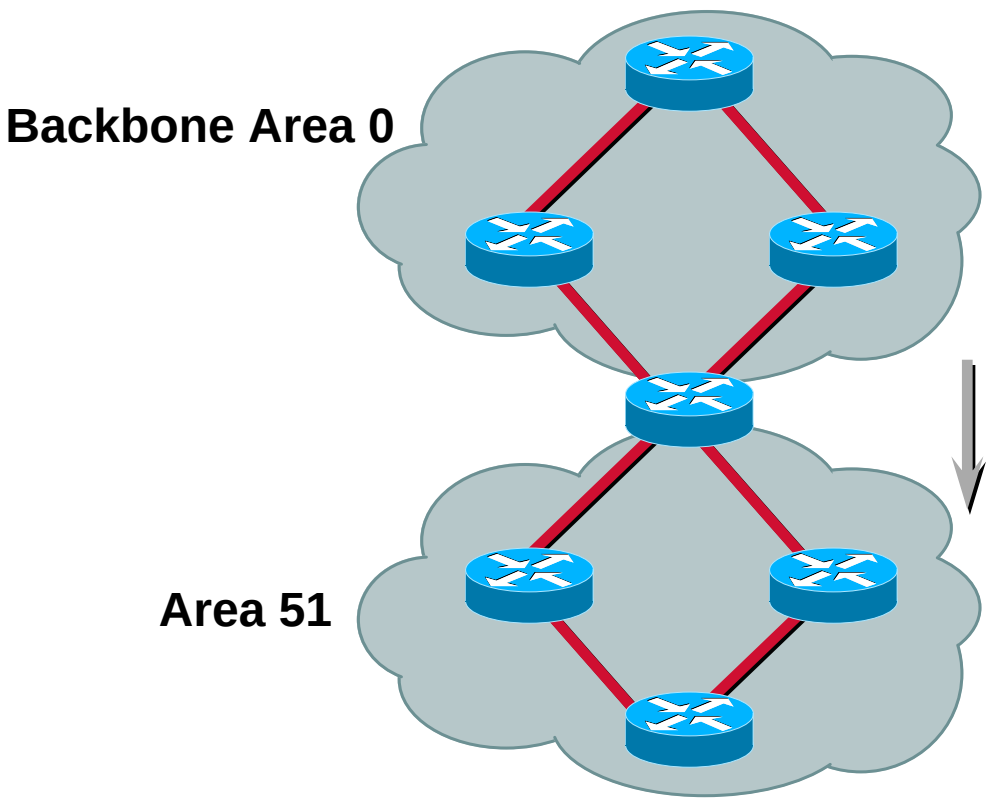

Type 3 and 5 LSA Filtering

A Totally Stubby Area Receives

Default Link 0.0.0.0

- **Use this for stable—scalable internetworks**

Not So Stubby Areas (NSSA)

Cisco.com

- **Benefits of stub area, but ASBR is allowed**
- **New type external LSA (type 7)**
  - **Type 7 LSAs flooded throughout the area**
  - **No type 5 external LSAs in the area**
  - **Type 7 LSAs will be converted into type 5 LSAs when flooded into area 0 by ABRs**
- **Filtering and summaries allowed at ABRs**

NSSA

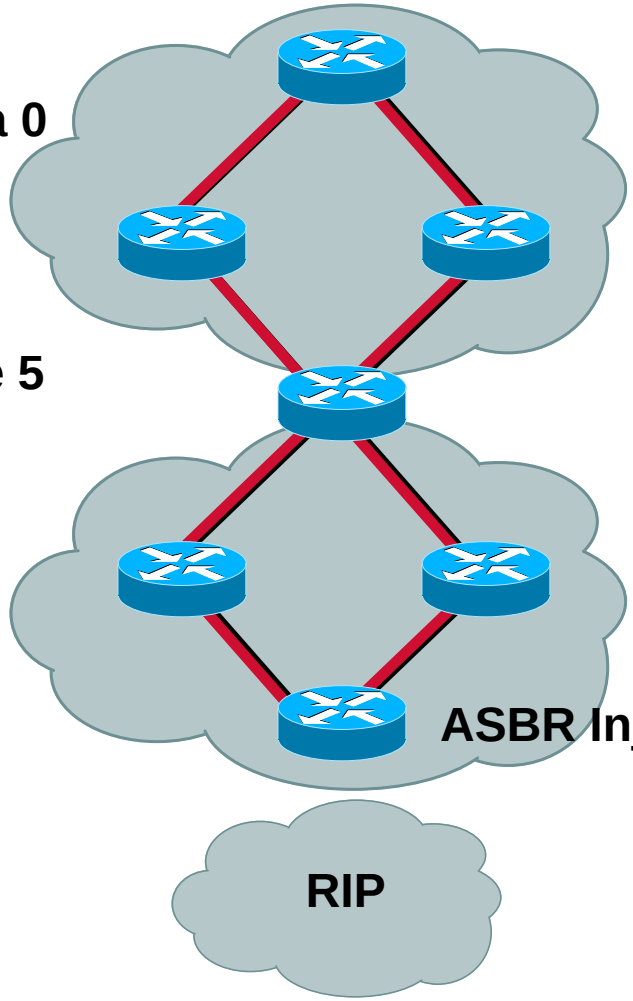

An NSSA Area Receives the Following

- Summary LSAs
- Specific LSAs
- Default External 0.0.0.0

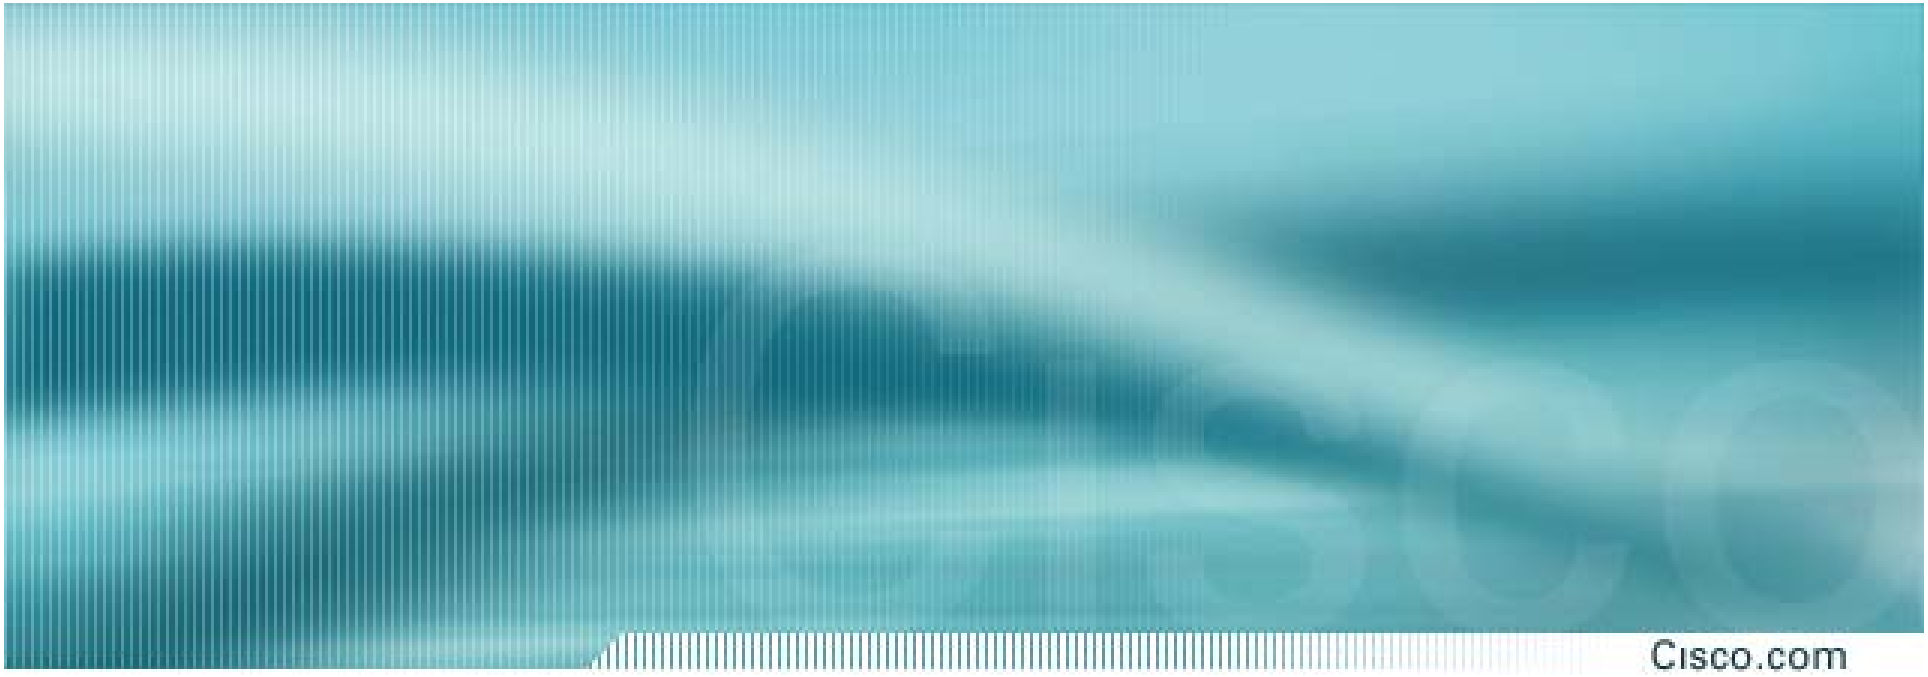

OSPF Media Options

OSPF Media Options

- **Multi-access media**
- **Point-to-point**
- **Non-Broadcast Multi-Access (NBMA)**
- **Demand circuits (11.2)**

Multi-Access Media

- Gig/Fast/Ethernet, FDDI, Token Ring
- Multicast
- DR and BDR

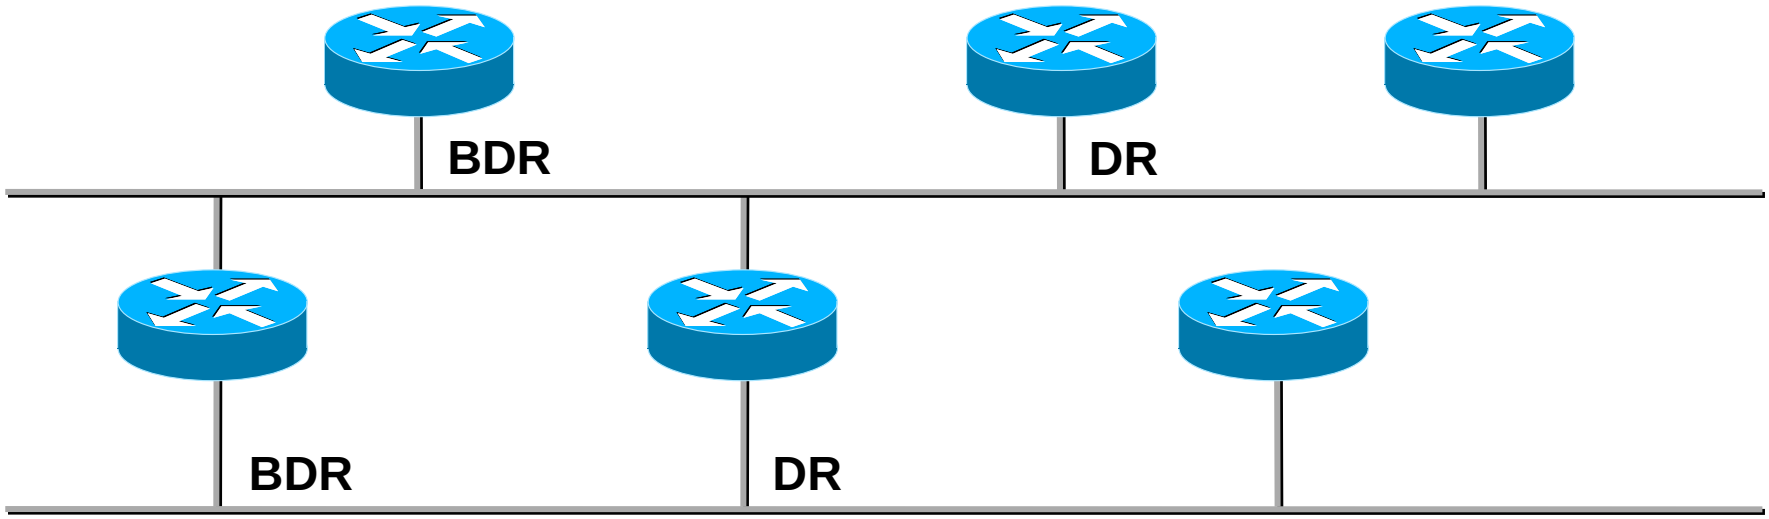

Designated Routers

Cisco.com

- Reduce OSPF traffic on multiaccess links
- Store and distribute neighbors LSDBs
- Backup DR for redundancy
- OSPF priority used in DR selection
- Range 1–255 default 1, 0 for noncandidate

Point-to-Point Media

- **Serial links**
- **Multicast used**
- **No DR or BDR**

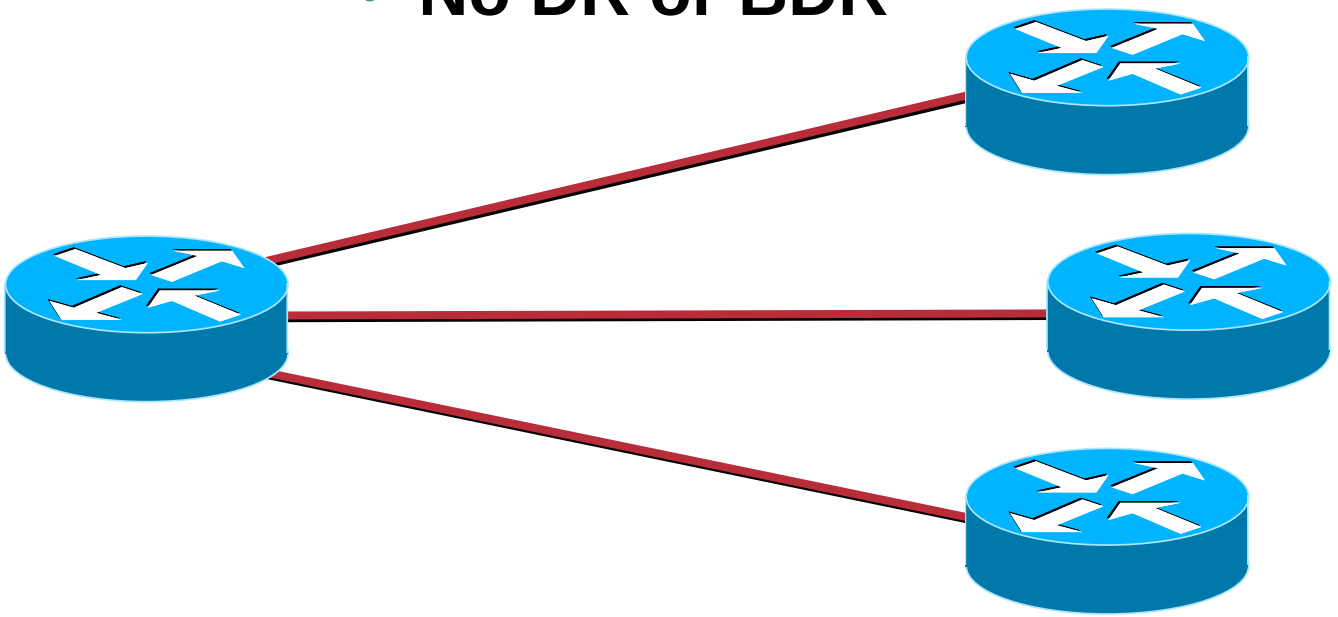

Non-Broadcast Multi-Access Media (NBMA)

- **Frame Relay (multipoint), X.25**

  Several possibilities: point-to-point, broadcast, point-to-multipoint, or nonbroadcast

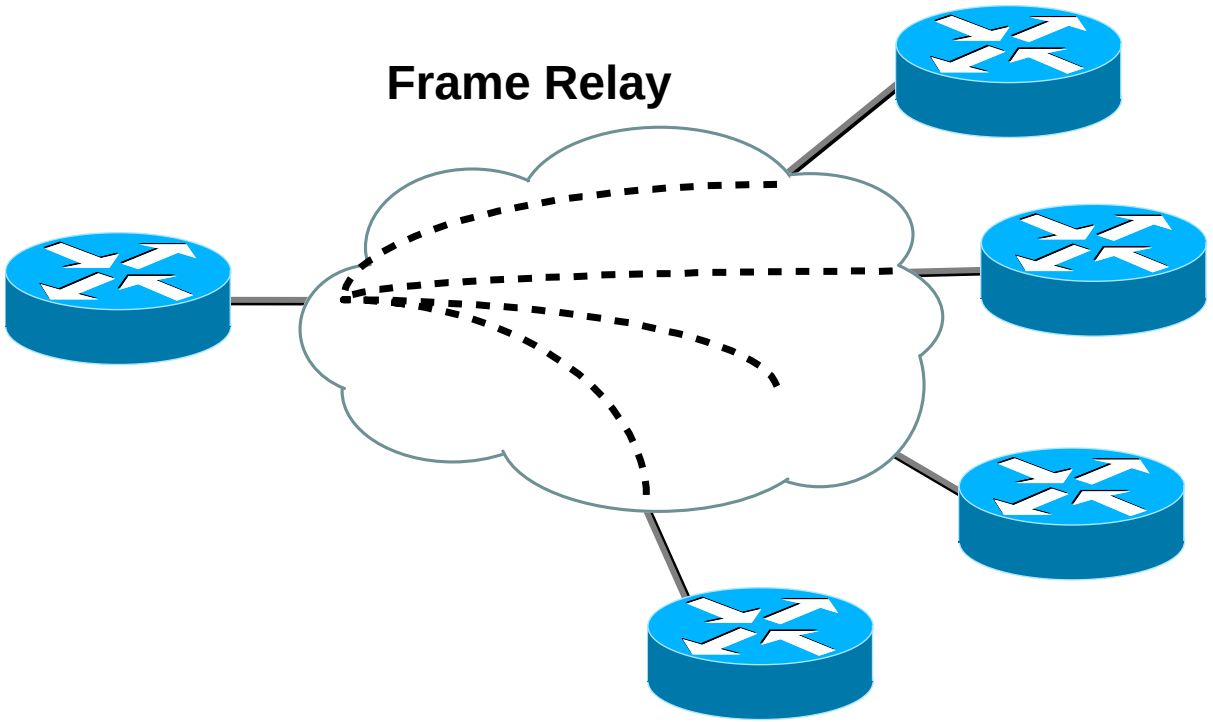

Dealing with NBMA

- **Point-to-point model**

  Benefits: individual costs can be configured; can be simple, treated like standard point-to-point links

  Drawbacks: complex to configure if the NBMA network is big or redundant; wastes address space

Dealing with NBMA

- **Broadcast model**

  Benefits: simple to configure; treated like a multi-access network

  Drawbacks: must maintain an L2 full-mesh at all times; one one metric for all VCs

Dealing with NBMA

- **NBMA model:**

  Benefits: only one IP subnet used

  Drawbacks: complex to configure and scale; need to manually configure each neighbor

Dealing with NBMA

- **Point-to-multipoint model:**

  Benefits: simple to configure; no neighbor configuration (unless you want individual costs); no requirement for a full mesh at L2

  Drawbacks: compared to other choices—none

- **This is the recommended method of dealing with NBMA networks**

OSPF Demand Circuits

- **OSPF demand circuit**
  - **Additional option in LSAs: do not age bit**
  - **Suppresses hellos exchange**
  - **Suppresses DB synchronization**
- **All new LSA still have to be transmitted in the area**

Agenda

Cisco.com

- Pre-Requisite
- Link State Protocol Refresher
- **Fundamental Deployment and Analysis**
- **OSPF Configuration Examples**
- **ISIS Configuration Examples**

Market Segments

Cisco.com

- **Market Segments**
  - **a) Service Providers**
  - **b) Enterprise**
    - **Manufacturing**
    - **Retail**

Service Providers

- SP networks are divided into pops
- Transit routing information is carried via BGP
- IGP is used to carry next hop only
- Optimal path to the next hop is critical

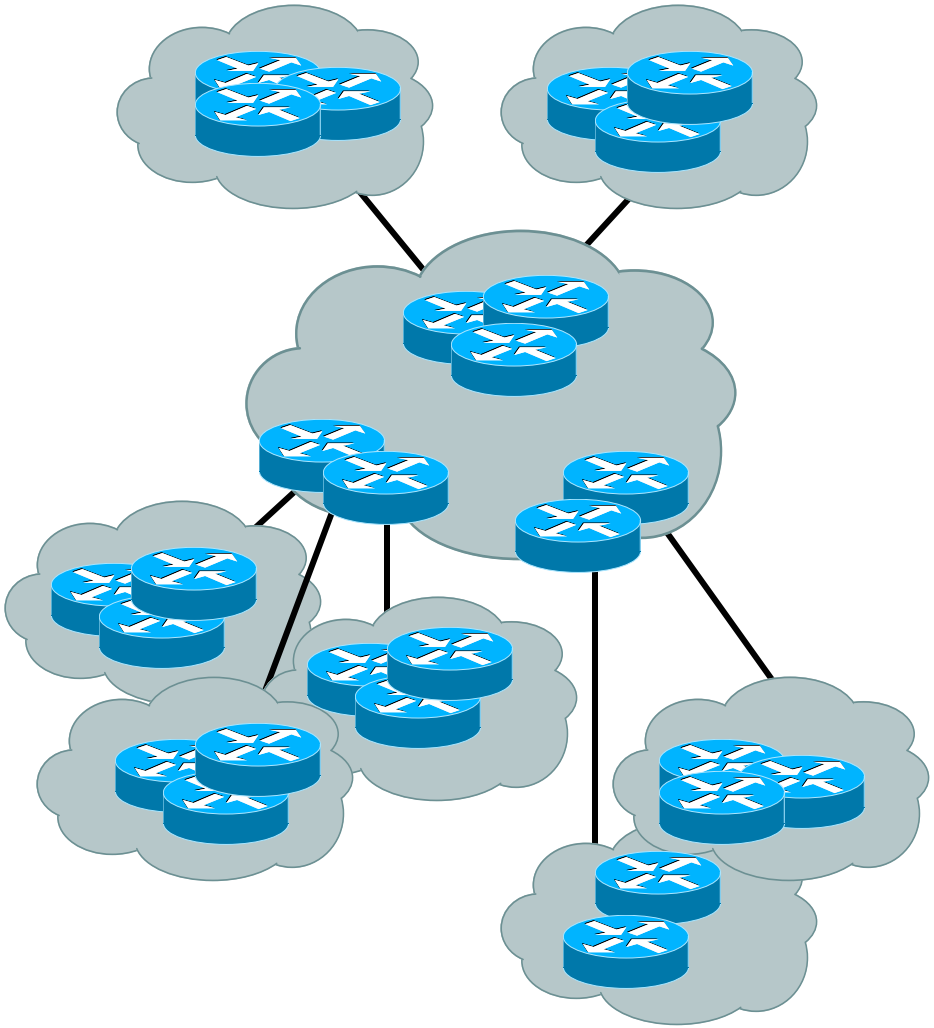

Enterprise Retails

- IGP scalability requirement on retail side is on Hub and Spoke side
- OSPF and ISIS were designed keeping the Service Provider infrastructure in mind
- Lacks clear vision towards hub and spoke architecture
- Acquisitions brings more topological restrictions
- Distance Vector are better choice for e.g., EIGRP, RIPv2, ODR
- BGP?

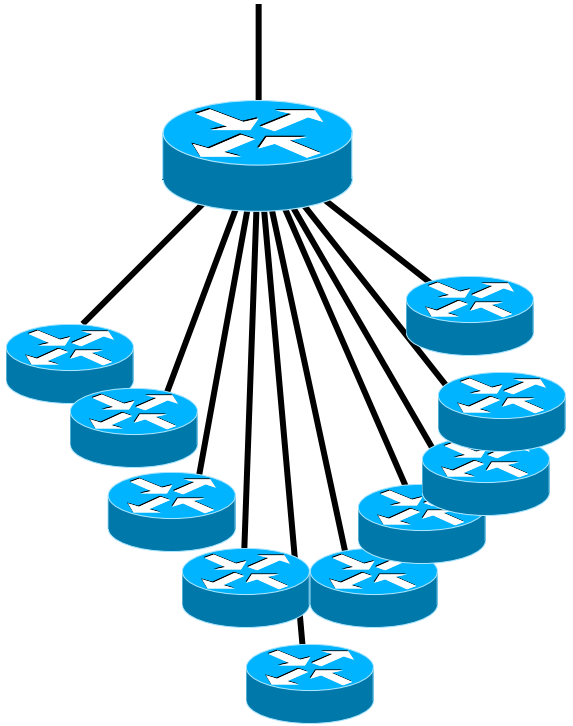

Enterprise Manufacturing

- **Closet type infrastructure for host communications**
- **Large number of end devices connections**
- **Campus networks**

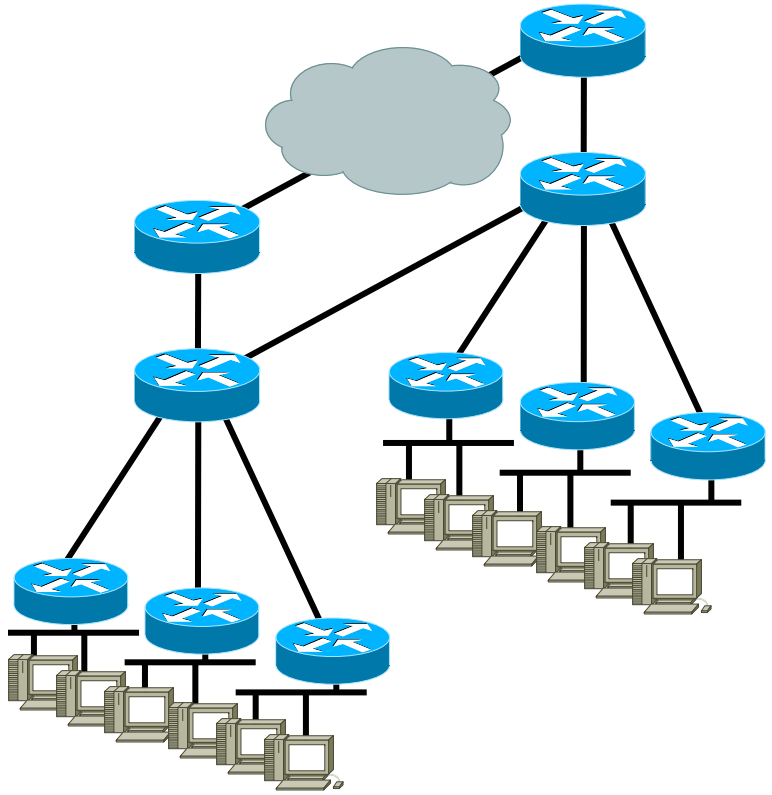

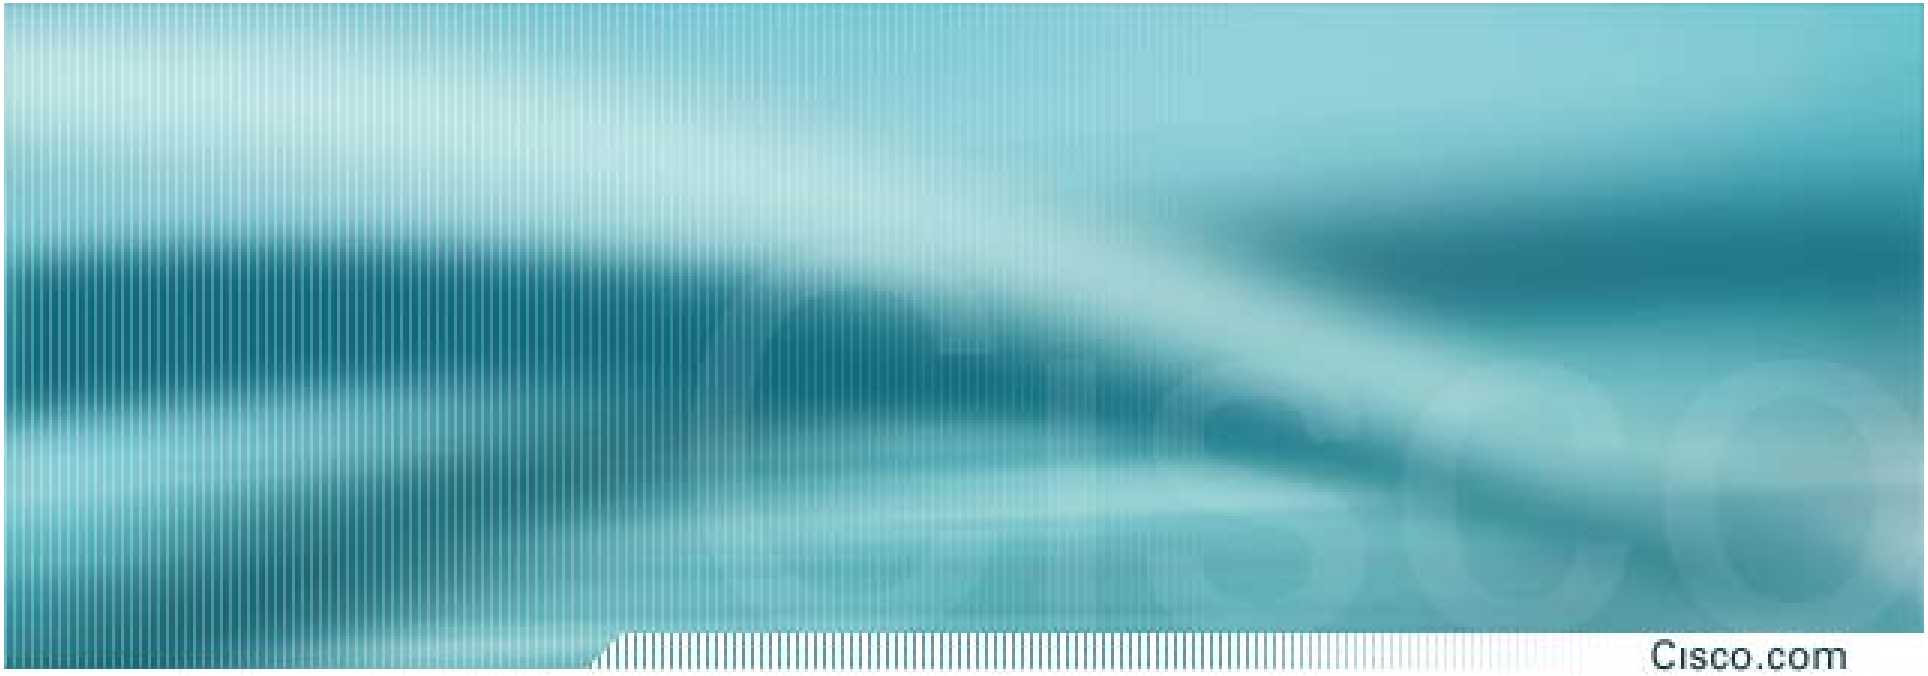

Service Providers Deployments

SP Deployment Characteristics

- SPs should have only one instance of IGP running throughout network (exceptions are there)
- BGP carries external reachability
- IGP carries only next-hop (loopbacks are better for e.g., next-hop-self)

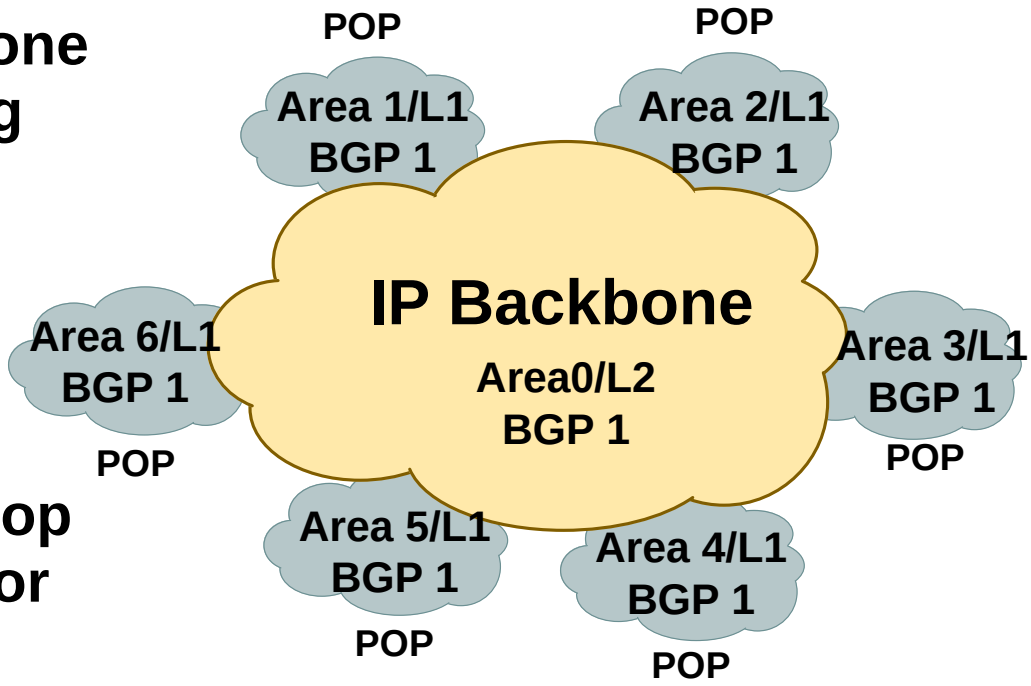

SP Architecture

- Major routing information is 110K via BGP
- Largest known IGP routing table is ~6–7K
- Total of 117K
- 6K/117K ~ 5% of IGP routes in an ISP network
- A very small factor but has a huge impact on network convergence!

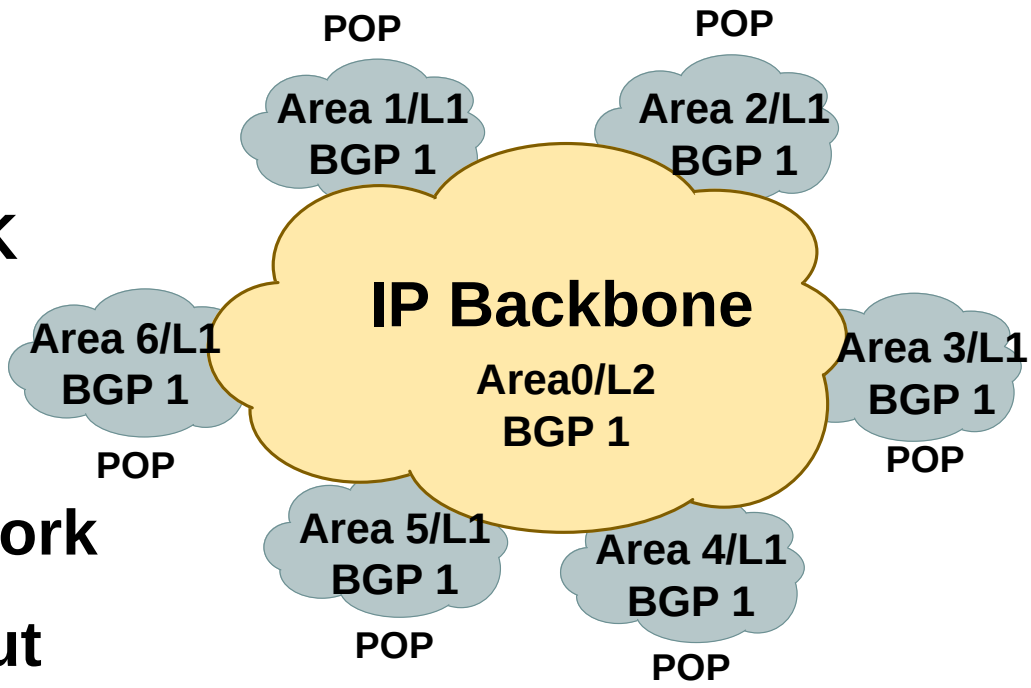

SP Architecture

- You can reduce the IGP size from 6K to approx the number of routers in your network
- This will bring really fast convergence
- Optimized where you must and summarize where you can
- Stops unnecessary flapping

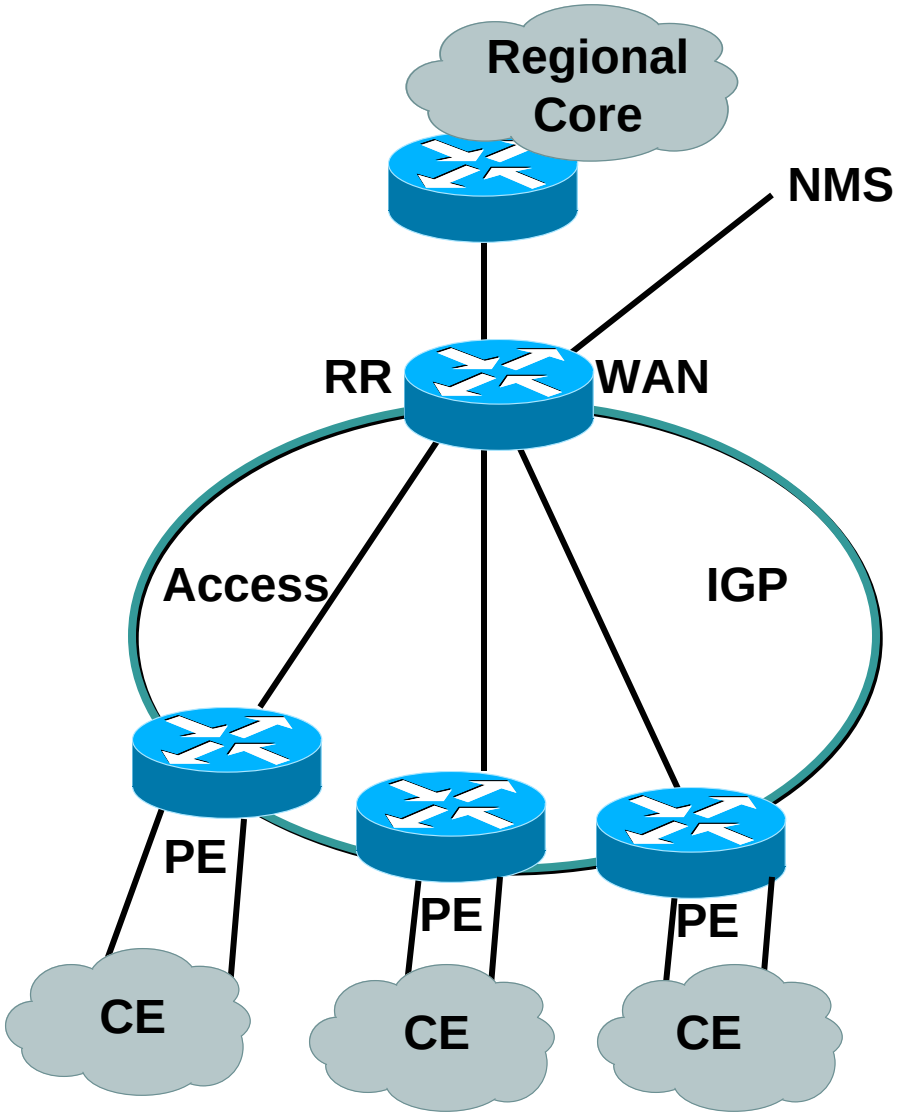

SP Architecture

- The link between PE-CE needs to be known for management purpose
- BGP next-hop-self should be done on all access routers—unless PE-CE are on shared media (rare case)
- This will cut down the size of the IGP
- For PE-CE link do redistributed connected in BGP
- These connected subnets should ONLY be sent through RR to NMS for management purpose; this can be done through BGP communities

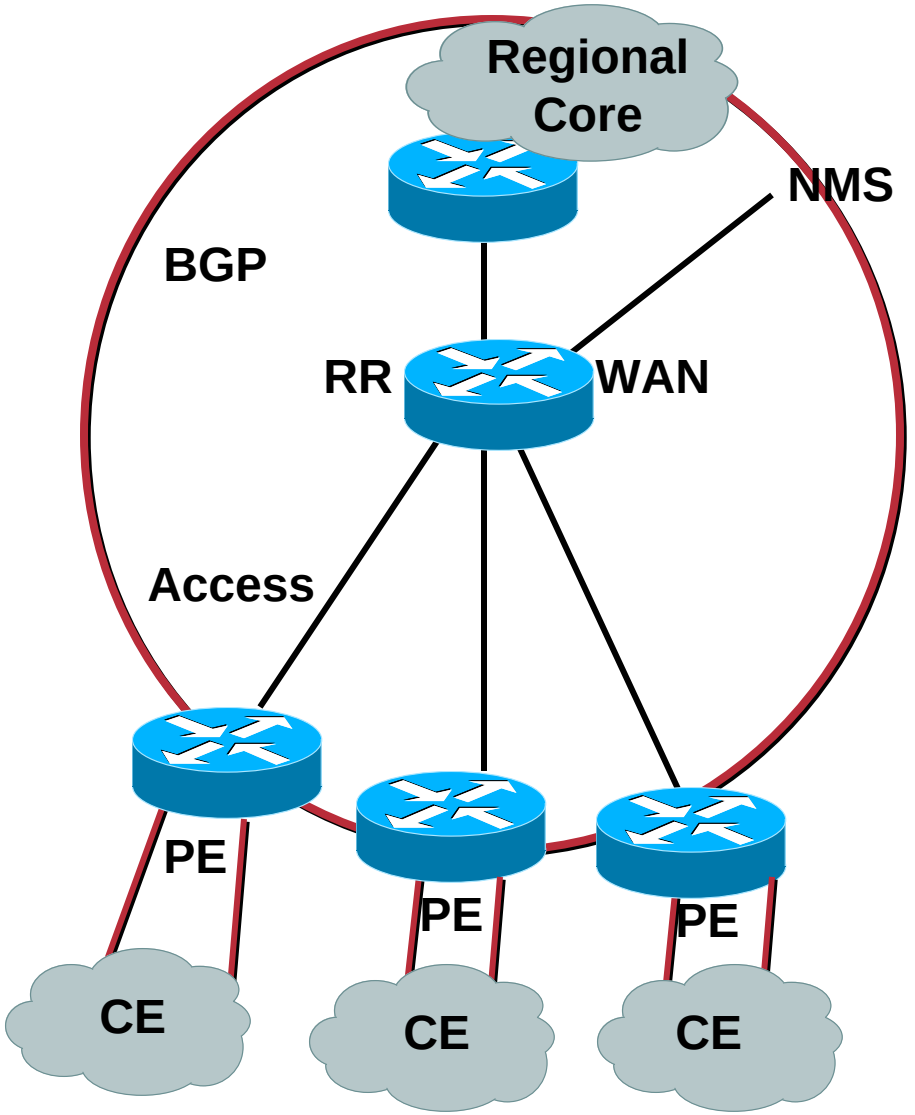

SP Architecture

- **Where do we define area boundaries? WAN routers can be L1L2 in ISIS or ABR in case of OSPF**
- **Hide the pop infrastructure from your core**
- **Traffic engineering if needed can be done in core from WAN routers**

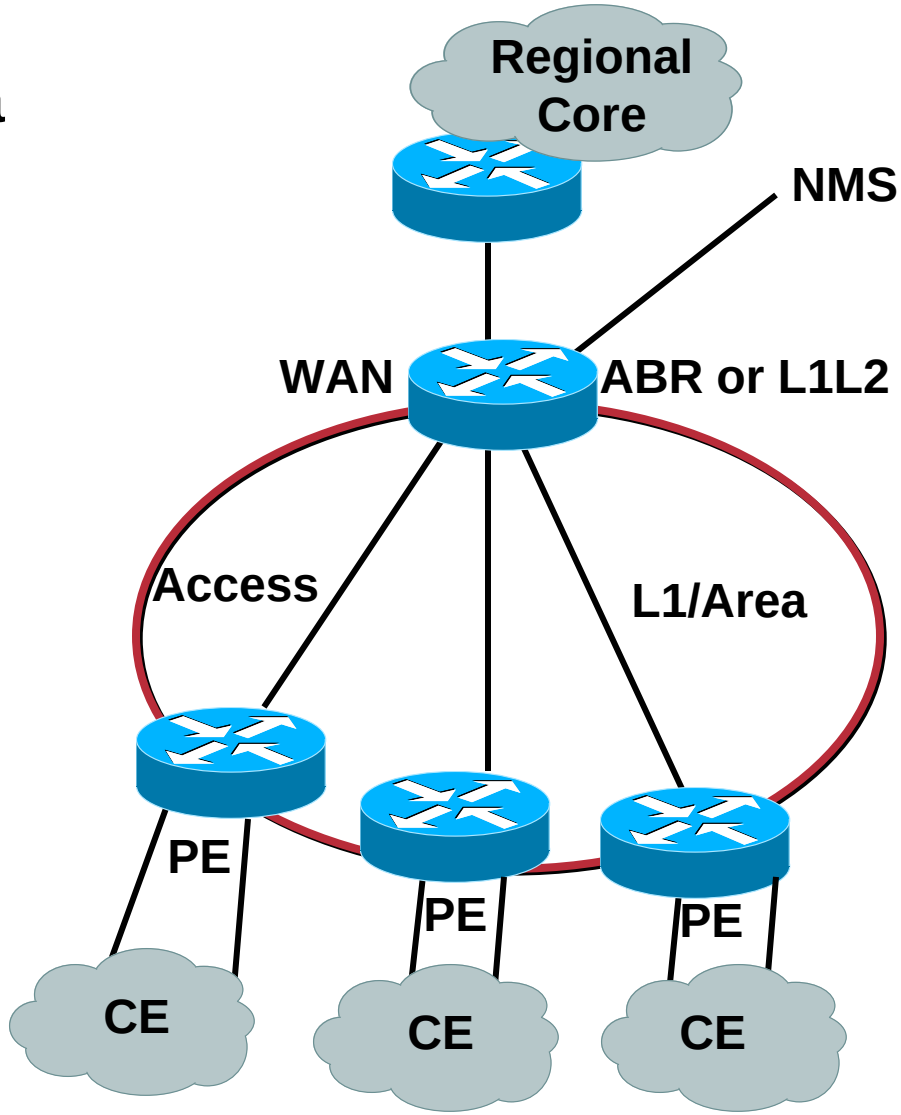

SP Architecture

- Physical address between ABR and PE should be in a contiguous blocks
- These physical links should be filtered via Type 3 filtering from area 0 into other areas
- Why? To reduce the size of the routing table within each pop
- Every area will carry only loopback addresses for all routers
- Only NMS station will keep track of those physical links
- PE device will not carry other Pop's PE's physical address in the routing table

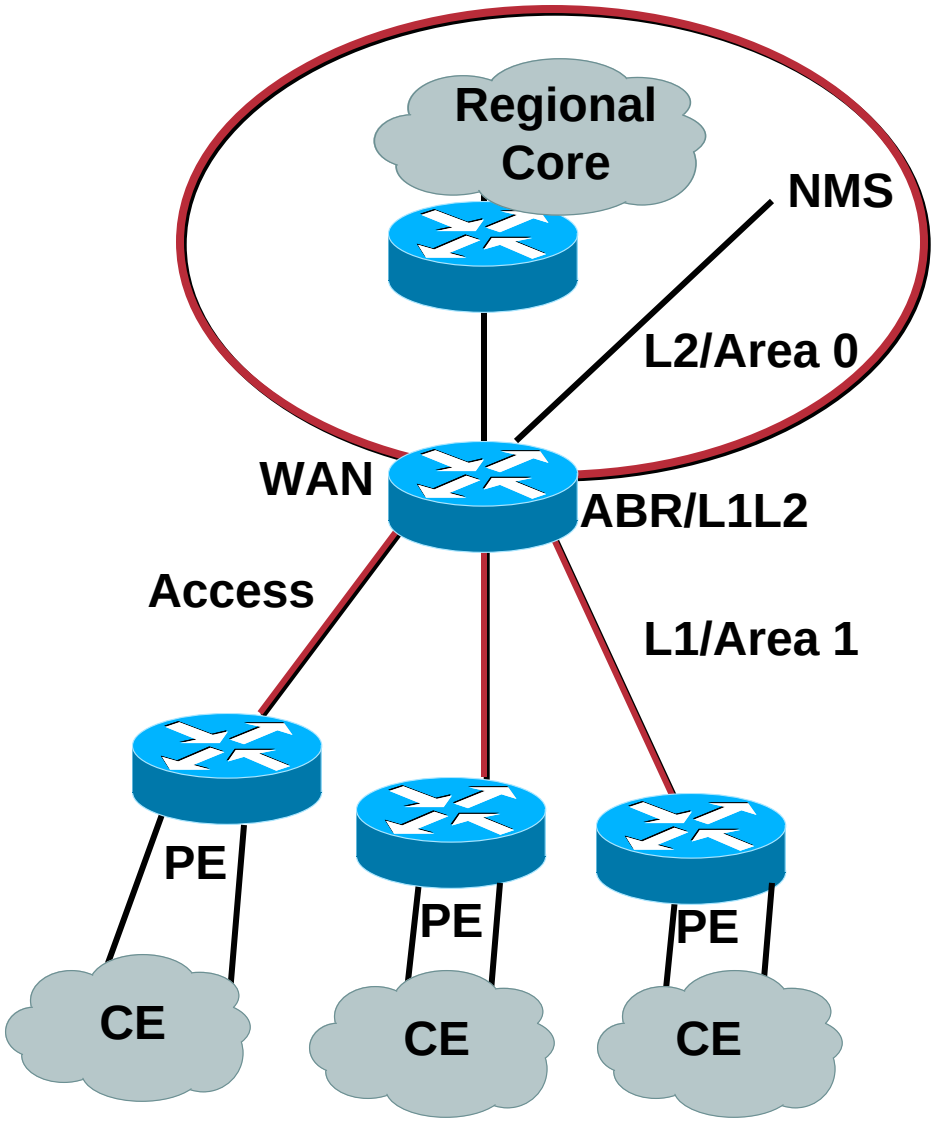

SP Architecture

- Area 0 will contain all the routes
- This is the most intelligent form of routing and also there will not be too many routes in IGP
- If there are 500 pops and every pop contains 4 routers; then instead of having 6K routes you will only have 2K
- This is scalable and hack proof network!

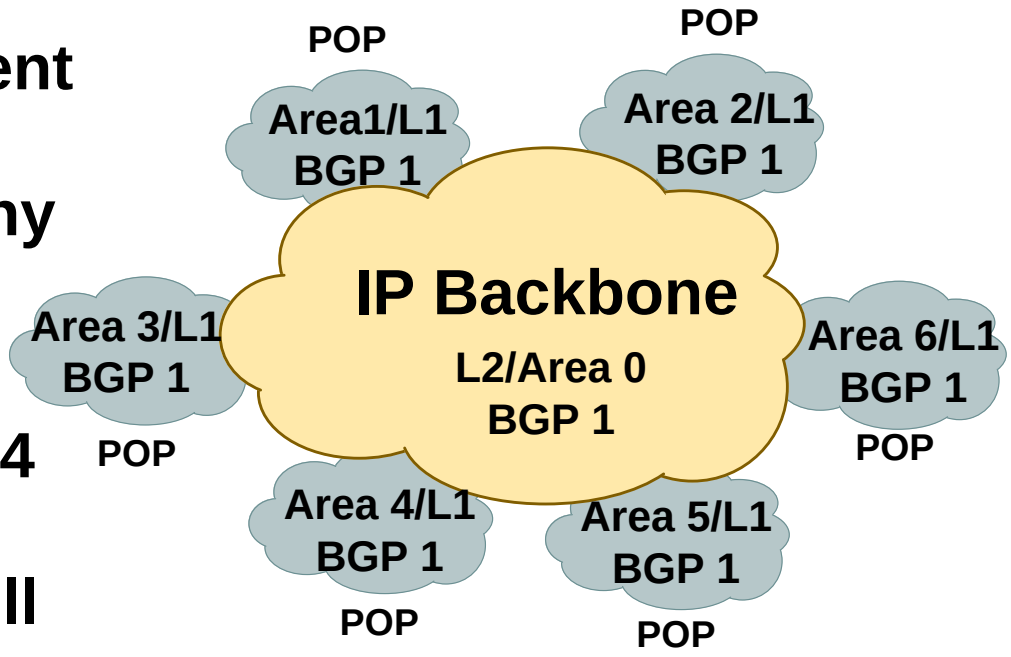

Dealing with Redistribution

- **Don’t do it!**
- **If you’re an SP you shouldn’t be carrying external information in your IGP**
- **Let BGP take care of external reachability**
- **Use OSPF or ISIS to carry only next-hop information—i.e., loopbacks**

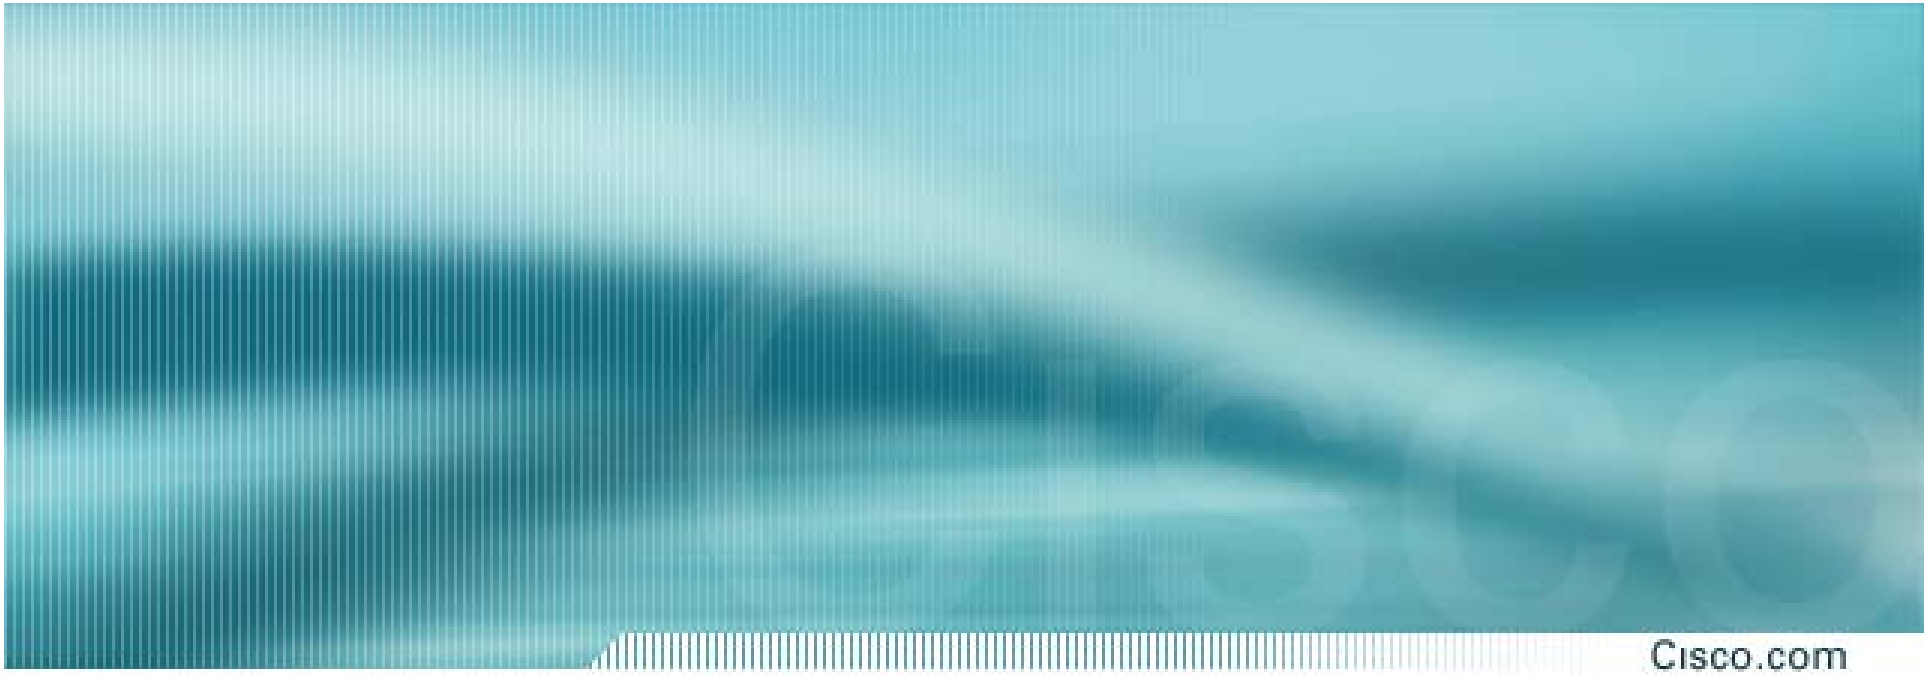

Enterprise Deployments

Enterprise Retail

- OSPF and ISIS were not very good choice for hub and spokes
- EIGRP, ODR, RIPv2 and BGP are better choice here
- Enterprise BGP is not complicated
- You do not need to play with lot of attributes

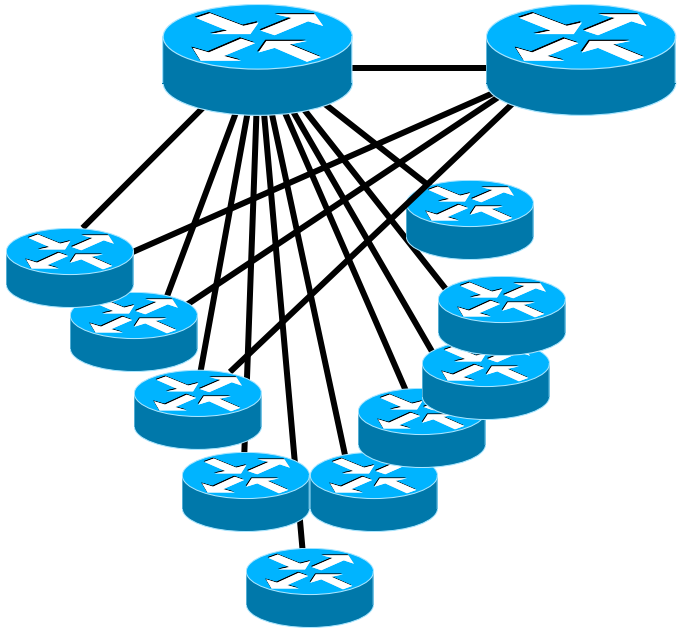

Enterprise Retail

- The link between 2 hub routes should be equal to the number of areas
- Summarization of areas will require specific routing information between the ABR's
- This is to avoid suboptimal routing
- As you grow the number of areas, you will grow the number of VLAN/PVC's
- This is protocol limitation

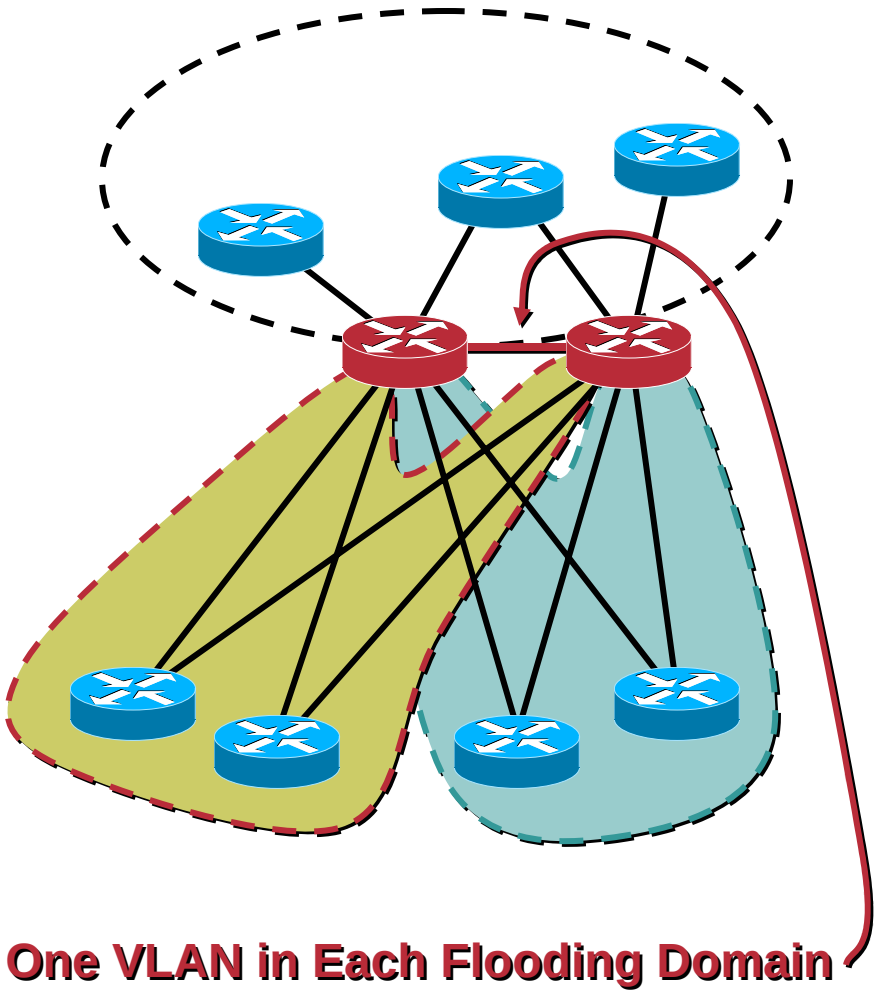

Enterprise Retail

- Spoke router in one area will contains route for the other spoke router since they are in the same area
- Acquisitions and merger will create another sets of problem
- Rearrangement of topology required if the area or the router limitation has been reached
- Very difficult to preserve the protocol's hierarchical design

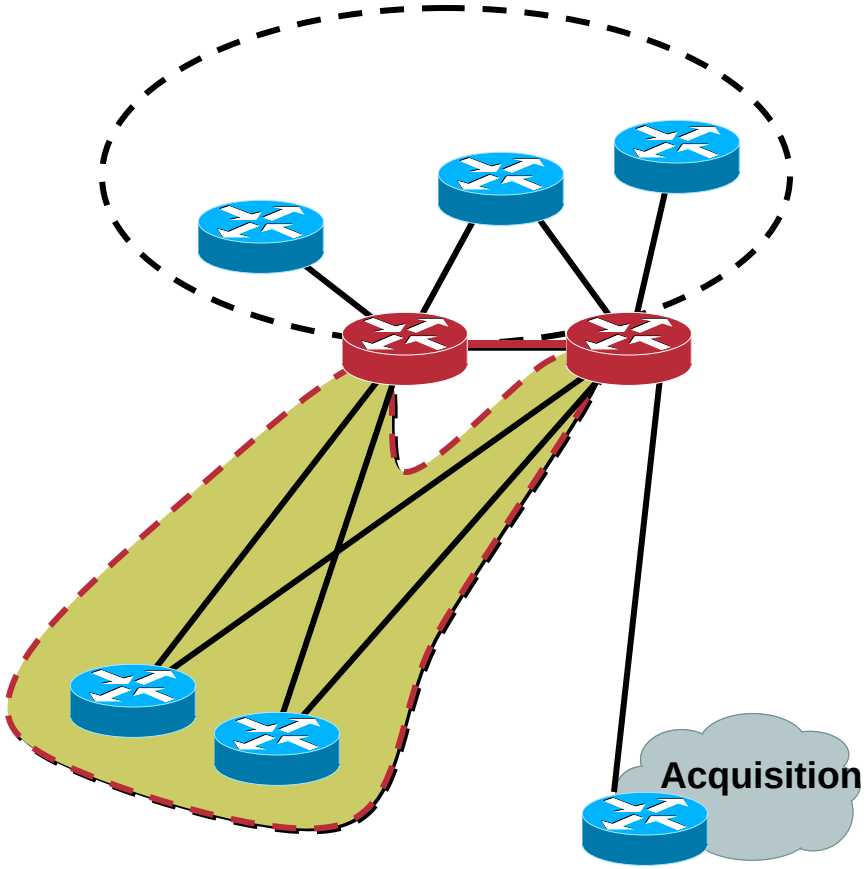

Enterprise Manufacturing

- Can have multiple ‘islands’ of IGPs
- Islands tied together by a BGP core
- May be a requirement for redistribution

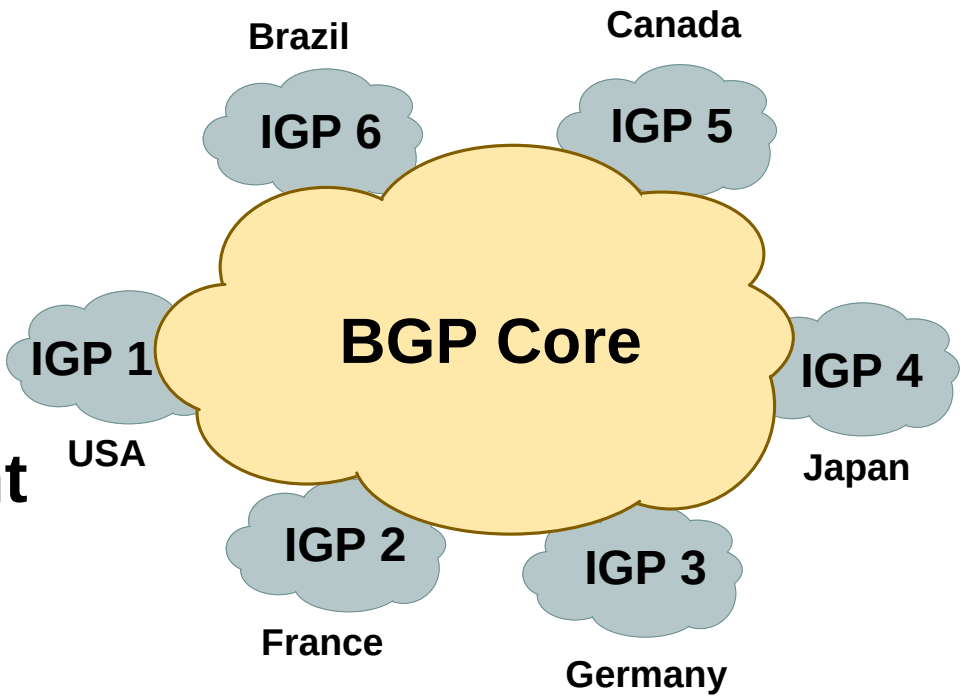

Dealing with Redistribution

- **The number of redistribution boundaries should be kept to a minimum**

  Why?

  Because you have better things in life to do besides; build access lists

- **When redistributing try to place the DR as close to the ASBR as possible to minimize flooding**

Dealing with Redistribution

- **Be aware of metric requirements going from one protocol to another**

  RIP metric is a value from 1–16

  OSPF Metric is from 1–65535

- **Include a redistribution default metric command as a protection**

```
router ospf 1
    network 130.93.0.0 0.0.255.255 area 0.0.0.0
    redistribute rip metric 1 subnets
```

Dealing with Redistribution

- **Redistribute only what is absolutely necessary**
  - **Don’t redistribute full Internet routes into OSPF**
- **Default route into BGP core; let the core worry about final destination**

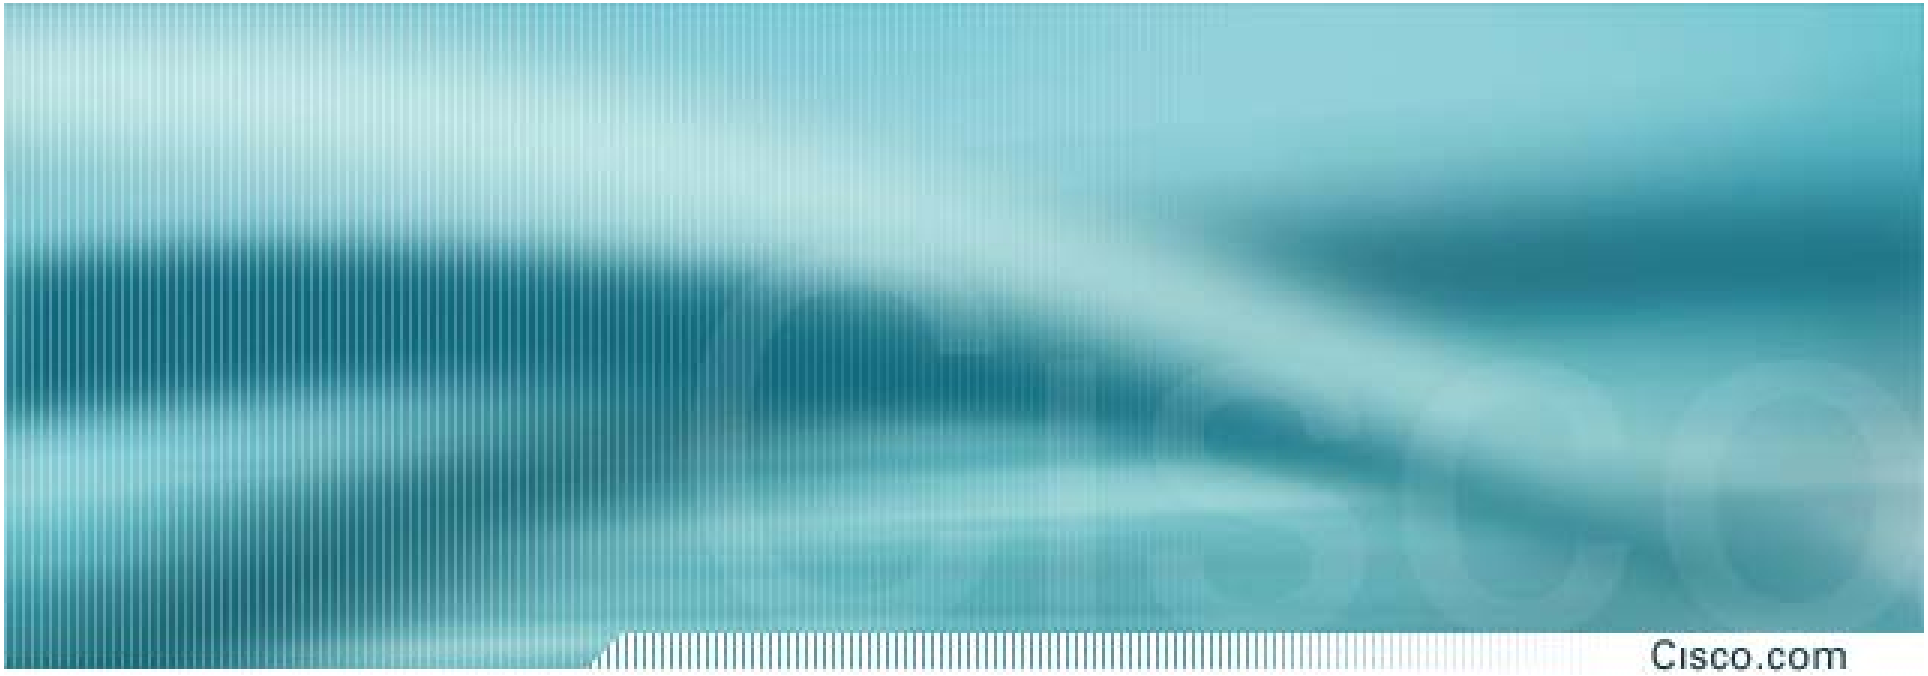

Design Practices

Accidental Redistribution Prevention (OSPF)

- Areas should be defined as stub to prevent accidental redistribution of eBGP into OSPF
- Type 3 LSA filtering should be used at ABR's and only routers' loopbacks should be allowed to leak into other areas
- Loopback should be in private address space to make LSA type 3 filtering easier; for e.g., 10.0.0.0/8
- iBGP routes can not be redistributed into IGP by default
- NMS resides in area 0 here

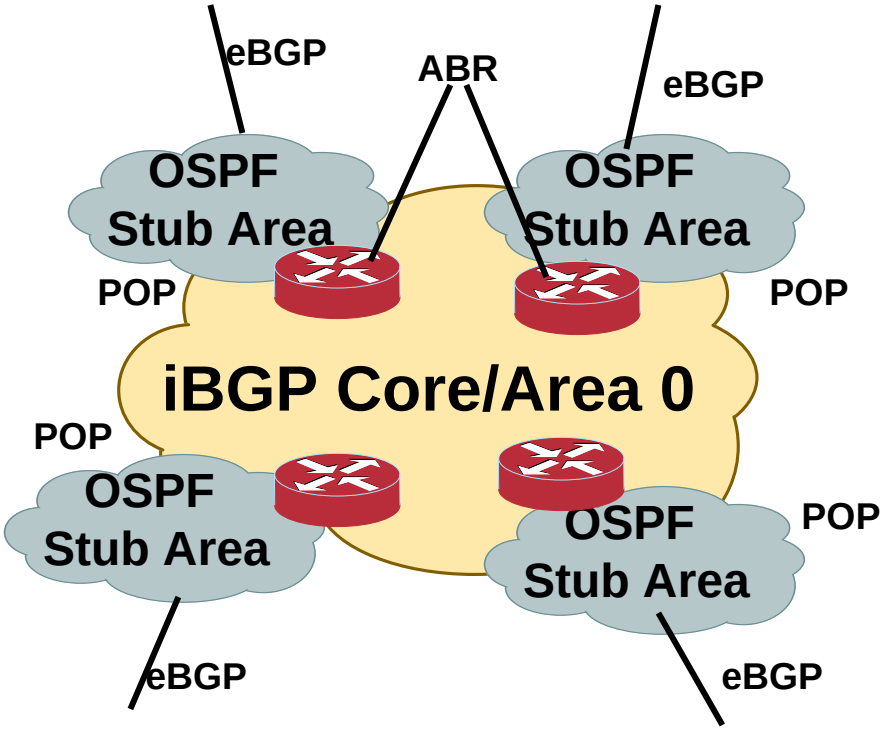

Accidental Redistribution Prevention (ISIS)

- L1 by default can not import more than 30K routes if eBGP is redistributed into ISIS by accident
- To save your backbone use summarization and only allow router's loopback address and physical links
- Route leaking should be done on all L1L2 routers to control all pops' info from leaking into other pops (for e.g., only leak 10/8)
- NMS resides in the L2/backbone here

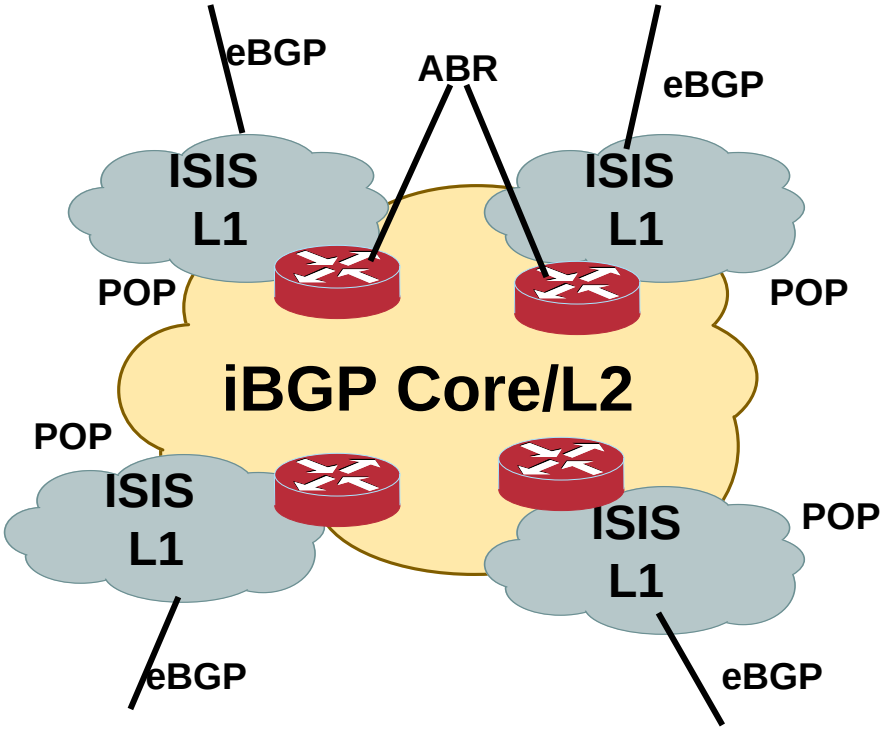

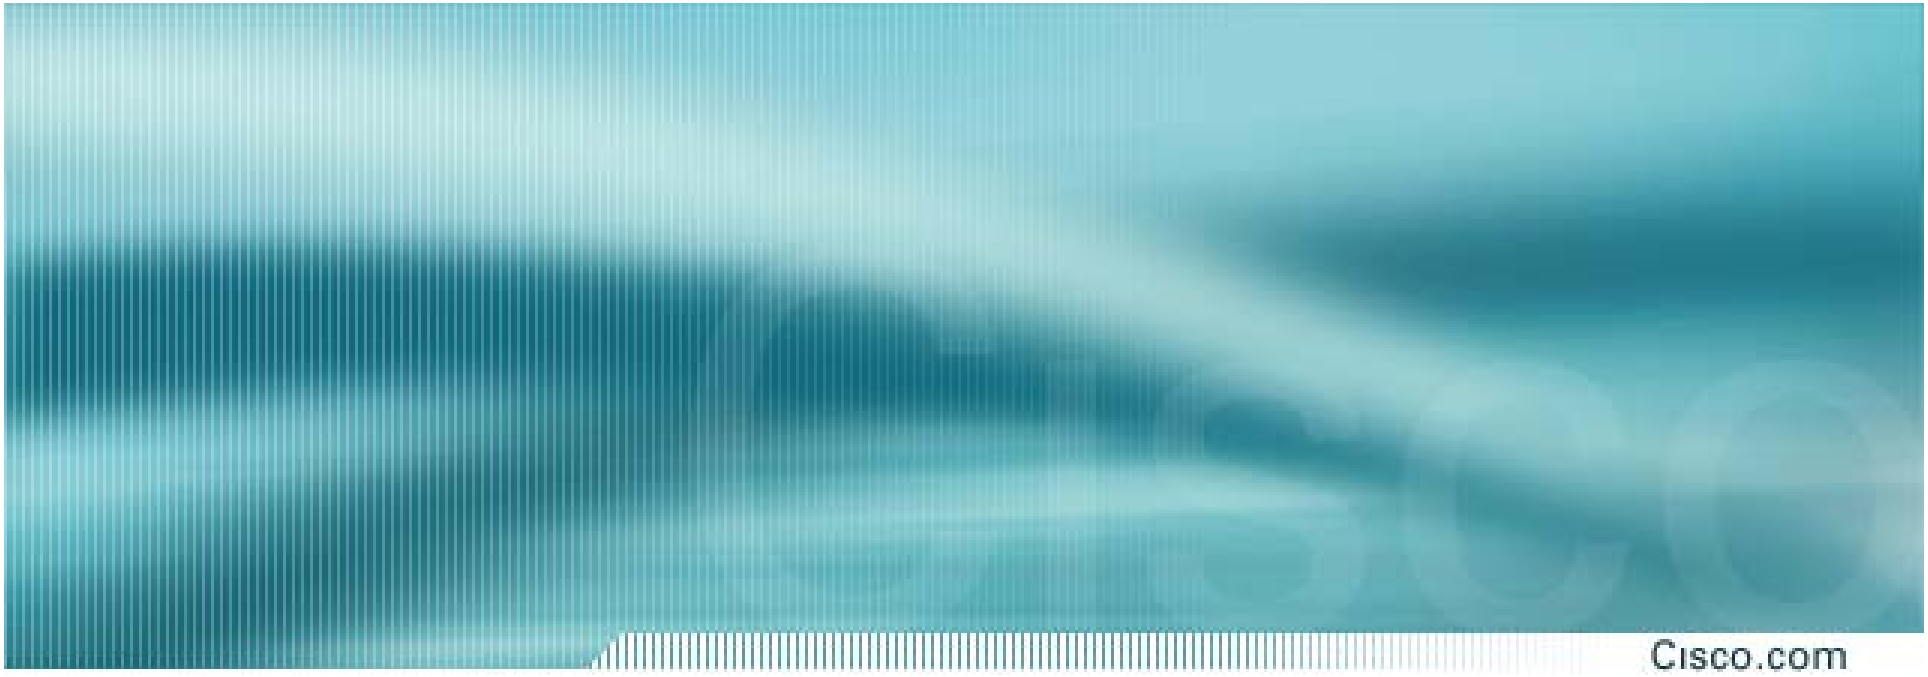

Dialup Design Tips

Dial Backup and OSPF

- **Dial-on-Demand Techniques:**
  - **Use virtual profiles**
  - **Virtual interface gets assigned to the area associated with the calling router**
  - **Area configuration and IP address for virtual interface are dynamic**

Dial Backup and OPSF

Cisco.com

Virtual Profiles

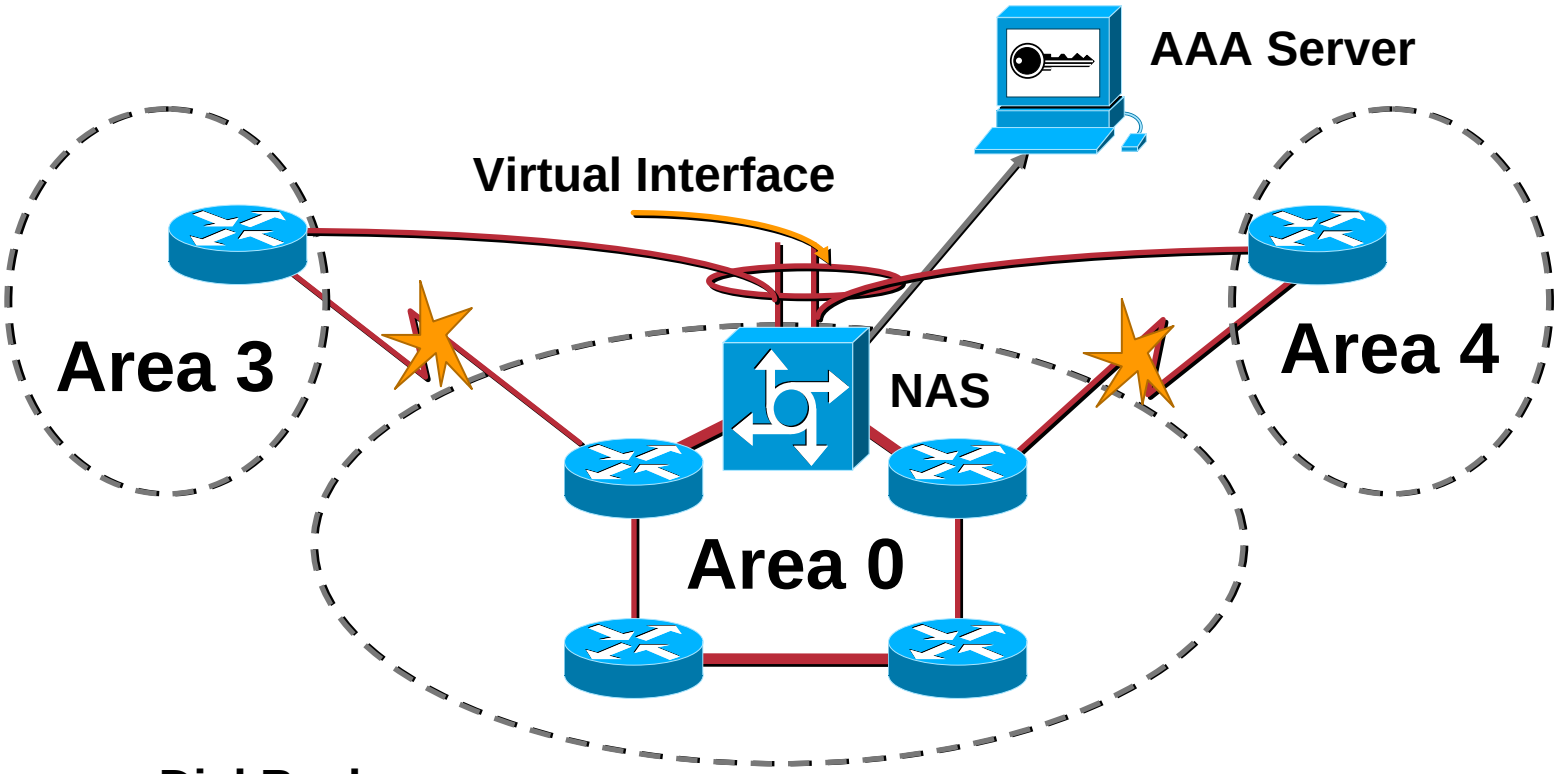

Dialup Design Practices

- **Two kinds of Pools can be defined on NAS: Static Pools and Distributed Pools**
- **Static Pool: address range remain within a single NAS—easier to manage from routing perspective**
- **Distributed Pool: address range may be distributed into multiple NAS's—hard to manage from a routing perspective**

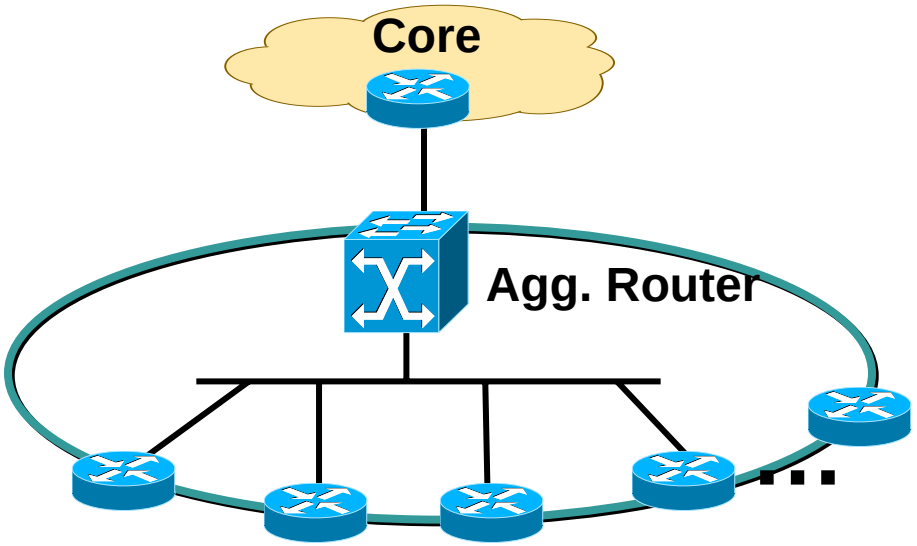

1000+ Routes Injected by Each NAS

Dialup Design with Static Pool Addresses

- **Three ways to propagate dialup routes from NAS:**
  - **Either Static route to pool address to null 0 with redistribute static on NAS or**
  - **Assign the pool add on a loopback on NAS with OSPF p2p network-type including loopback in an OSPF area or**
  - **Static route on ABR for the pool address pointing towards NAS (ASBRs)—this is a preferred method because summarization can be done at ABR**
- **Static pool do not require redistribute connected subnets on NAS**

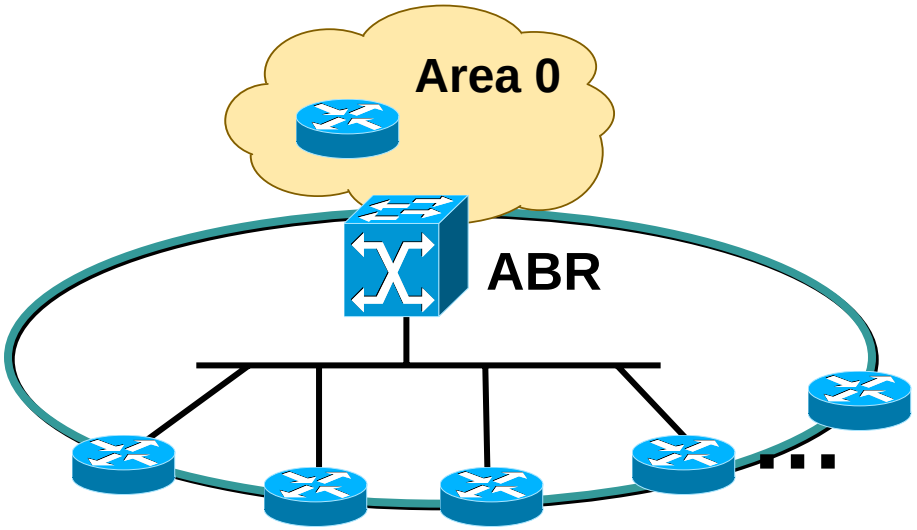

1000+ Routes Injected by Each NAS

Dialup Design with Distributed Pool Addresses

- **Distributed pool REQUIRES *redistribute connect subnets***
- **If pool is distributed, you can't summarize the pools at ABR because of *redistribute connected subnets* on NASs' unless its an NSSA, why?**
- **NSSA can summarize routes at ABR or ASBR**

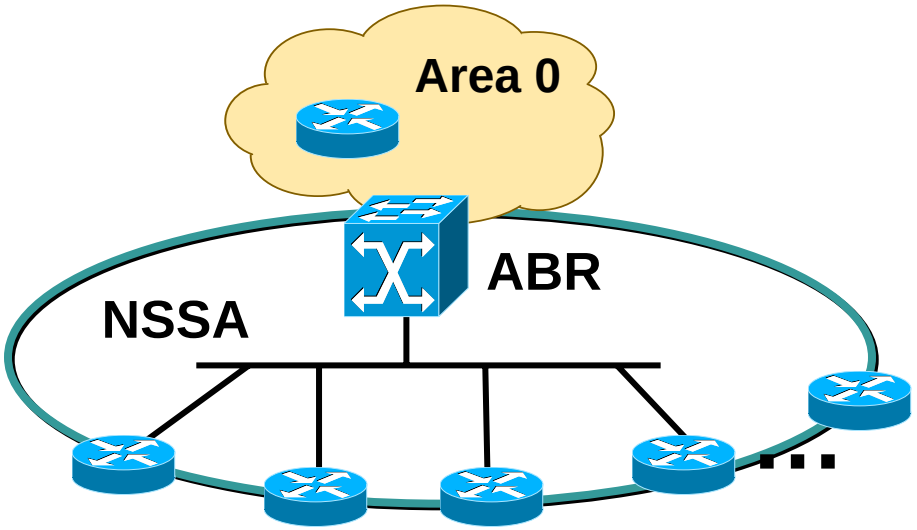

1000+ Routes Injected by Each NAS

Dialup Design Practices Scalability Issues

- If an area has too many routes injected by NAS then break it up in more than one area
- Area should be configured as NSSA for controlling type 5 at ABR level
- NSSA ABR can filter type 5 originated by NAS servers
- Configure totally NSSA so one area type 5 will not go into other areas

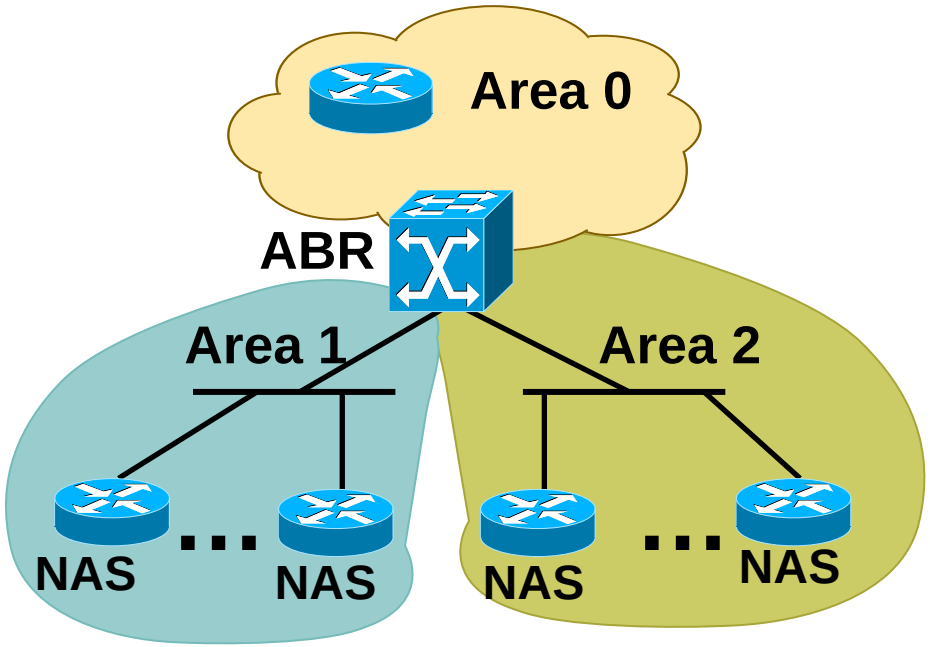

1000+ Routes Injected by Each NAS

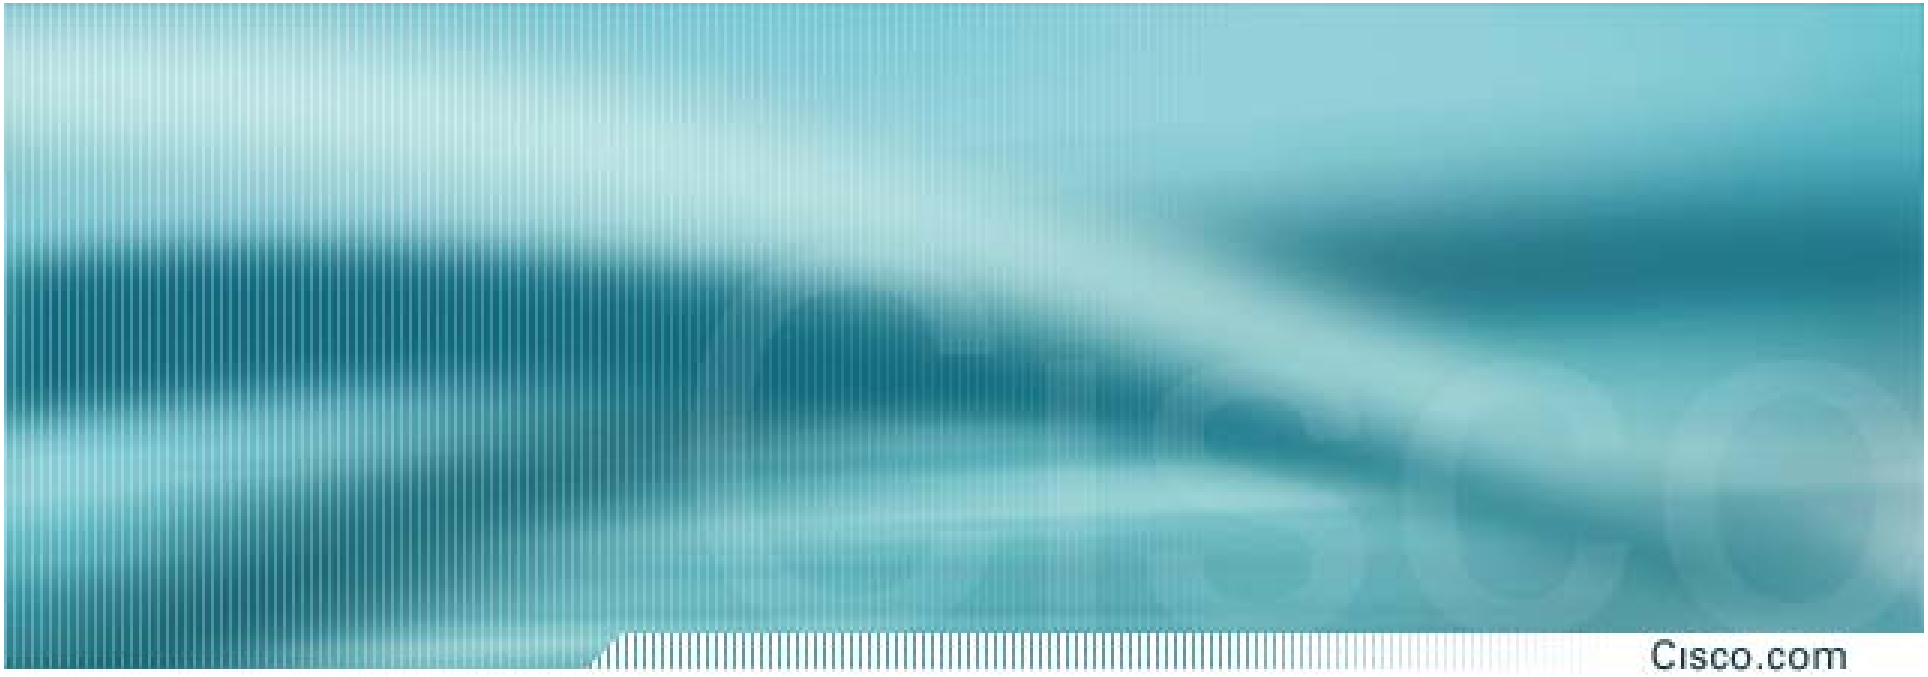

Migration Tips

Migration Approaches

Cisco.com

- **Technique 1: run only 1 IGP in any given part of the network; redistribute between other IGP and OSPF/ISIS**
- **Techniques 2: run two IGPs concurrently, giving OSPF/ISIS a higher admin distance and then cutting over**

Migration via Redistribution

- **Benefits: requires less router resources**
- **Drawbacks: becomes complex very quickly; requires careful redistribution with potentially lots of filters**
- **Not really a great choice for migration, but OK for small implementations**

Migration via Ships in the Night

- **Benefits: Cleaner process not requiring redistribution points in the network; easier to verify validity before making major change; easier to back out of**
- **Drawbacks: Requires more router resources, particularly RAM**
- **The recommended method of migrating to OSPF/ISIS; simpler, cleaner, safer, plus RAM is cheap these days**

Migrating

Cisco.com

- **Start by building out the L2/backbone area**
- **Migrate additional areas in incrementally**

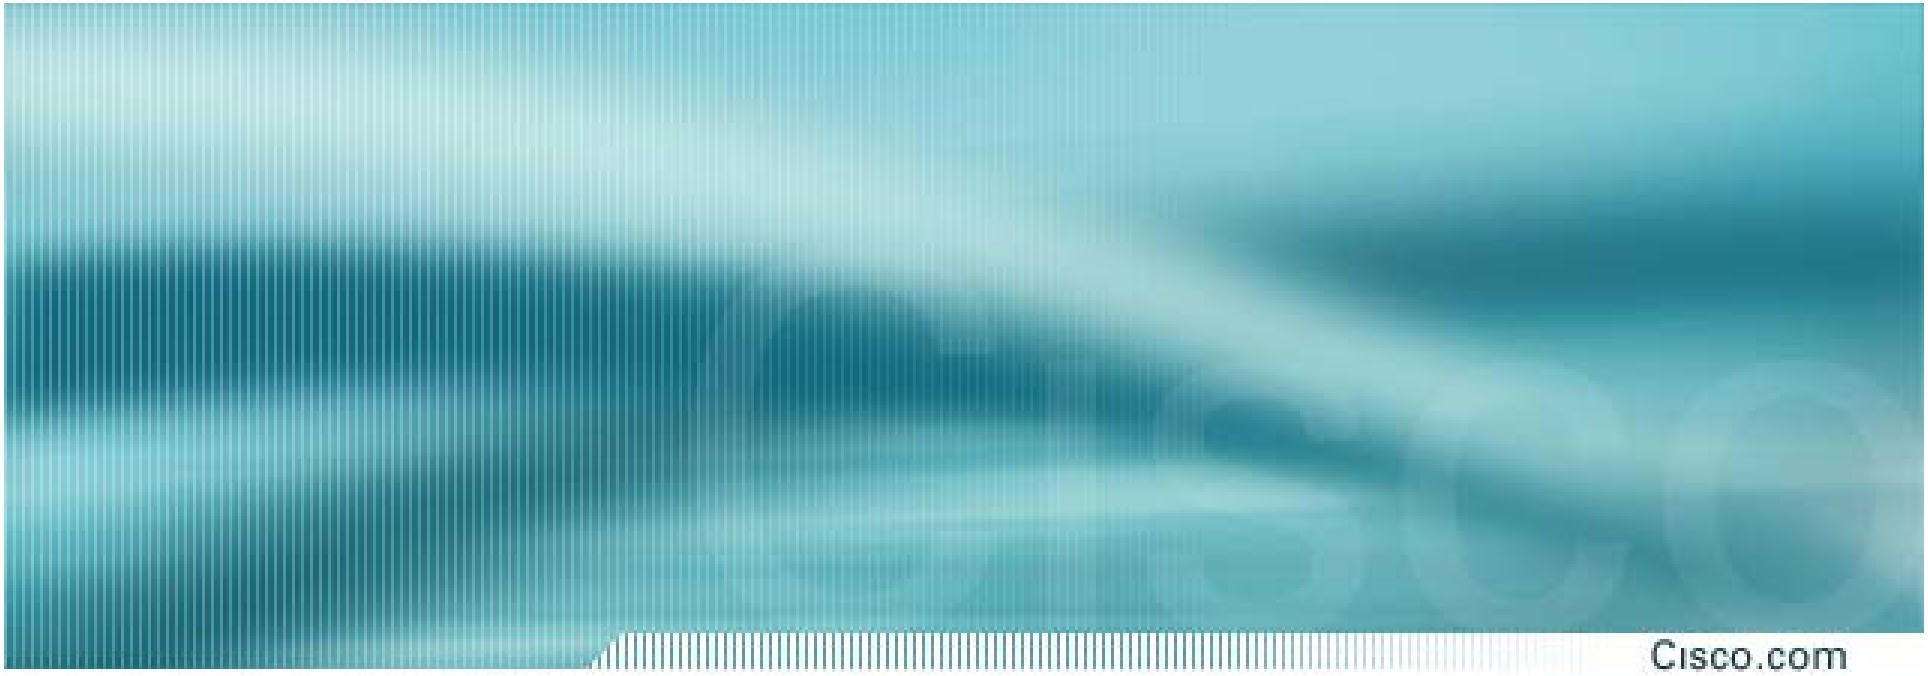

Protocol Validation

Protocol Validation

- **Do I have “too many” routers in an area?**

  It varies from network to network

- **Are the route ages appropriate?**

  If the protocol is periodic, routes should never be older than the update time

  If the protocol is update based, routes should get old

Protocol Validation (Cont.)

- **Is the number of routes stable?**
- **If you have the luxury, fail a link, allow the network to converge then restore the link and let the network converge**
  - **Are convergence times appropriate for the protocol**
  - **Remember to account for differentials in convergence time**

Protocol Validation (Cont.)

- **Do I have all the neighbors I should?**
- **Are routes getting old**

  show ip route

  Look at the age of the route

Agenda

Cisco.com

- Pre-Requisite
- Link State Protocol Refresher
- Fundamental Deployment and Analysis
- **OSPF Configuration Examples**
- **ISIS Configuration Examples**

Addressing

- Loopbacks are /32 out of 172.168.1.0
- WAN links are /30 out of 172.168.2.0
- Area 1 subnet: 172.168.6-7.0/24
- Area 2 subnet: 192.64.48-63.0/24
- Area 3 subnet: 64.78.134-170.0/24
- Area 4 subnet: 64.78.180-183/24
- Area 51 subnet: 173.45.12-26.0/24

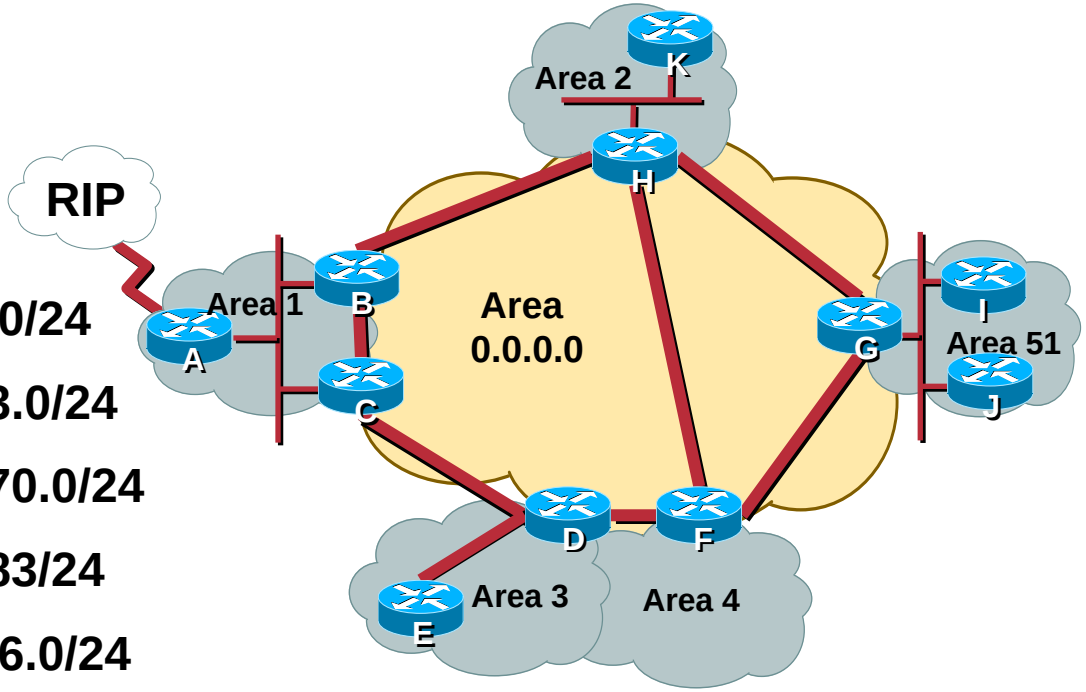

Design Plan

- Make Area 2 totally stubby
- Make Area 51 a stub
- Make Area 1 an NSSA
- All other areas are regular areas
- Summarize where possible
- Make router A the DR for the LAN in Area 1

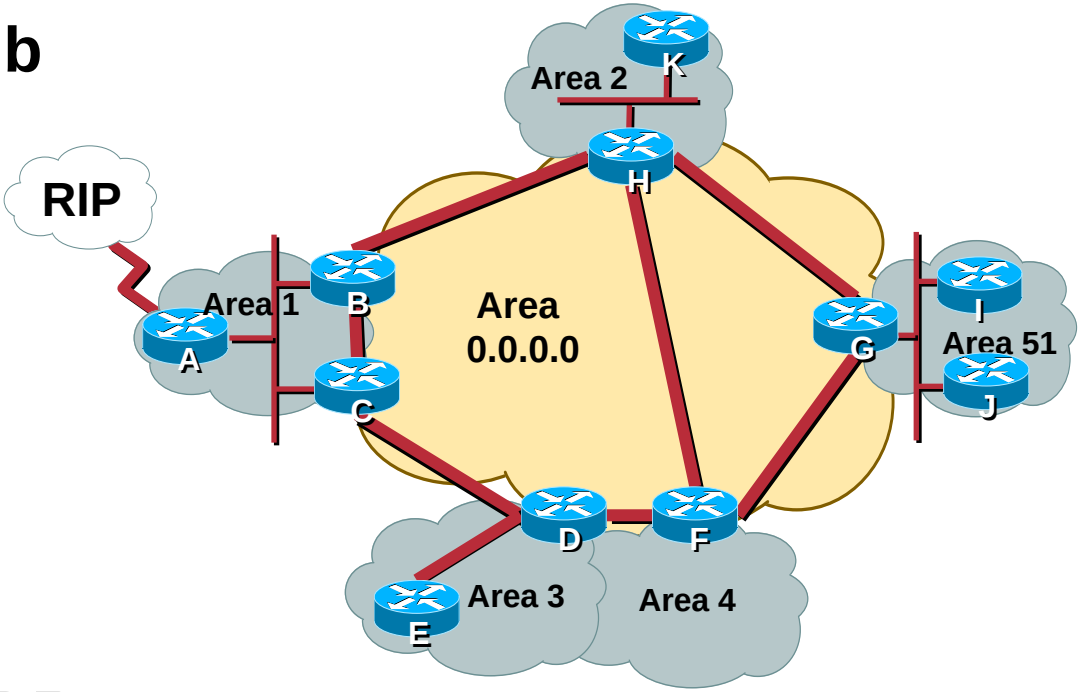

Example Configs—Area 1

Cisco.com

Router A

```
int lo0
 ip address 172.168.1.1 255.255.255.255
!
int se0
 ip address 192.145.100.235 255.255.255.252
!
int eth0
 ip address 172.168.6.1 255.255.255.0
 ip ospf priority 20
 ip ospf message-digest-key 1 md5 mykey
!
router ospf 1
 network 172.168.6.0 0.0.0.255 area 0.0.0.1
 redistribute rip subnets metric-type 2 metric 100
 area 0.0.0.1 authentication message-digest
 area 0.0.0.1 nssa
!
router rip
 network 192.145.100.0
 default metric 5
!
end
```

Router B

```
int lo0
 ip address 172.168.1.2 255.255.255.255
!
int se0
 ip address 172.168.2.5 255.255.255.252
 ip ospf authentication-key mykey
!
int se1
 ip address 172.168.2.13 255.255.255.252
 ip ospf authentication-key mykey
!
int eth0
 ip address 172.168.6.2 255.255.255.0
 ip ospf message-digest-key 1 md5 mykey
!
router ospf 1
 network 172.168.6.0 0.0.0.255 area 0.0.0.1
 network 172.168.2.0 0.0.0.255 area 0.0.0.0
 area 0.0.0.1 range 172.168.6.0 255.255.254.0
 area 0.0.0.0 range 172.168.2.0 0.0.0.255
 area 0.0.0.1 authentication message-digest
 area 0.0.0.0 authentication message-digest
 area 0.0.0.1 nssa
!
end
```

Example Configs—Area 2

Cisco.com

Router H

```
int lo0
 ip address 172.168.1.8 255.255.255.255
!
int se0
 ip address 172.168.2.6 255.255.255.252
 ip ospf authentication-key mykey
!
int se1
 ip address 172.168.2.9 255.255.255.252
 ip ospf authentication-key mykey
!
int se2
 ip address 172.168.2.17 255.255.255.252
 ip ospf authentication-key mykey
!
int eth0
 ip address 192.64.60.1 255.255.255.0
 ip ospf message-digest-key 1 md5 mykey
!
router ospf 1
 network 172.168.2.0 0.0.0.255 area 0.0.0.0
 network 192.64.60.0 0.0.0.255 area 0.0.0.2
 area 0.0.0.2 range 192.64.48.0 0.0.15.255
```

Router H (cont.)

```
area 0.0.0.2 stub no-summary
 area 0.0.0.0 authentication message-digest
 area 0.0.0.2 authentication message-digest
!
end
```

Router K

```
int lo0
 ip address 172.168.1.11 255.255.255.255
!
int eth0
 ip address 192.64.60.34 255.255.255.0
 ip ospf priority 20
 ip ospf message-digest-key 1 md5 mykey
!
router ospf 1
 network 192.64.60.0 area 0.0.0.2
 area 0.0.0.2 stub no-summary
 area 0.0.0.2 authentication message-digest
!
end
```

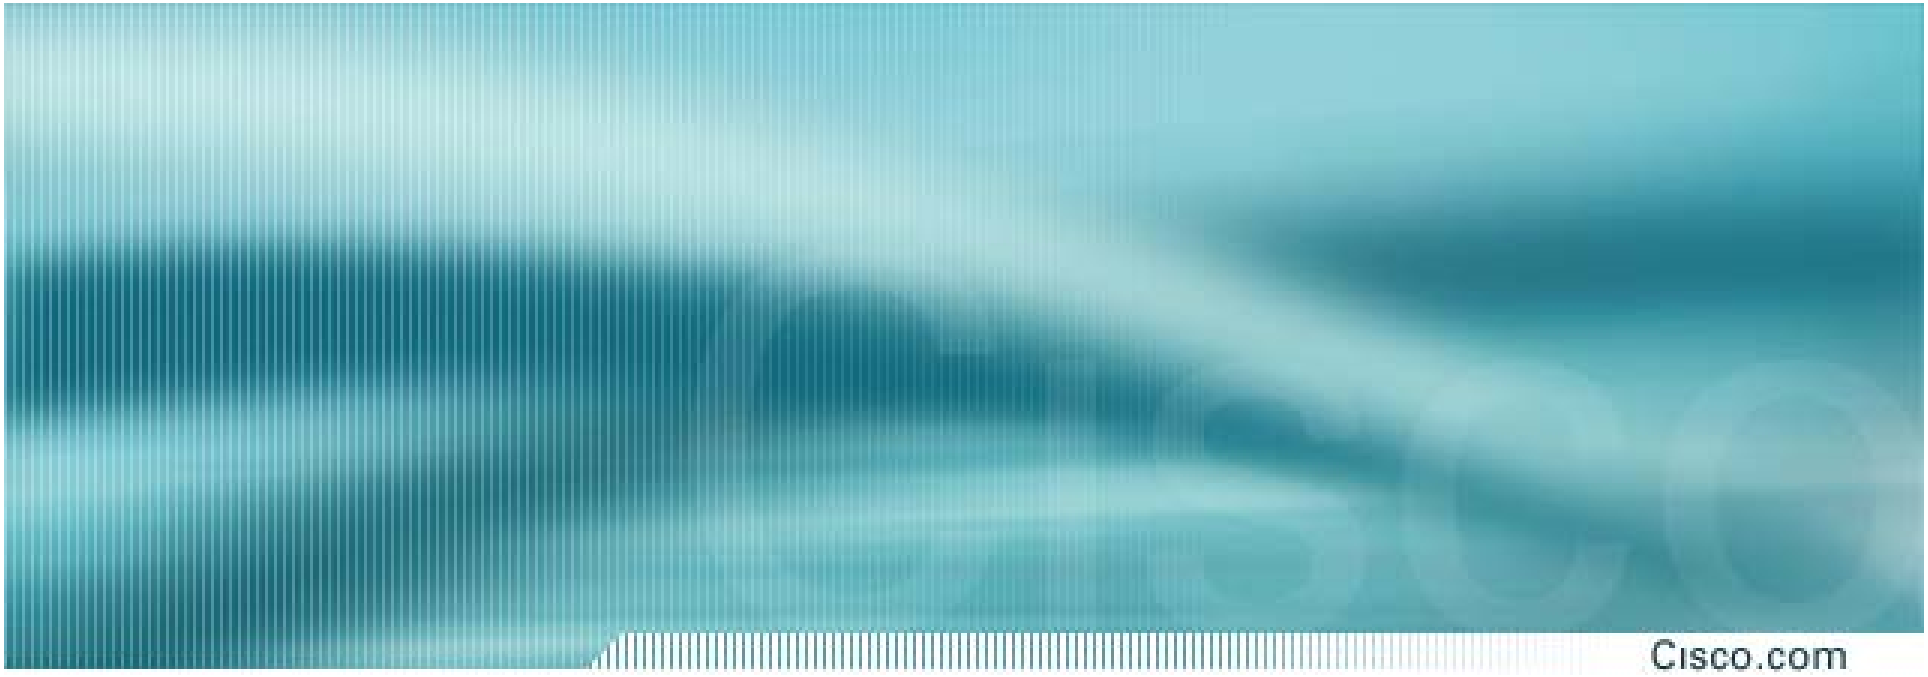

Tuning OSPF

Available “Nerd-Knobs”

- **Hello/Dead Timers**

  ip ospf hello-interval 3 (default 10)

  ip ospf dead-interval 15 (default is 4x hello)

  This allows for faster network awareness of a failure, and can result in faster reconvergence, but requires more router CPU and generates more overhead

- **LSA Pacing**

  timers lsa-group-pacing 300 (default 240)

  This is a great feature; allows grouping and pacing of LSA updates at configured interval; reduces overall network and router impact

Available “Nerd-Knobs”

- **DR/BDR Selection**

  ip ospf priority 100 (default 1)

  This feature should be in use in your OSPF network; forcibly set your DR and BDR per segment so that they are known; choose your most powerful, or most idle routers; try to keep the DR/BDR limited to pne segment each

- **OSPF Internal Timers**

  timers spf 2 8 (default is 5 and 10)

  Allows you to adjust SPF characteristics; first number sets wait time from topology change to SPF run; second is hold-down between SPF runs; BE CAREFUL WITH THIS COMMAND; if you’re not sure when to use it, it means you don’t need it; default is 95% effective

Available “Nerd-Knobs”

- **LSA filtering/interface blocking**

  Per interface:

  ip ospf database-filter all out (no options)

  Per neighbor:

  neighbor 1.1.1.1 database-filter all out (no options)

  OSPFs router will flood an LSA out all interfaces except the receiving one; LSA filtering can be useful in cases where such flooding unnecessary (i.e., NBMA networks), where the DR/BDR can handle flooding chores

  area <area-id> filter-list <acl>

  Filters out specific Type 3 LSAs at ABRs

- **Improper use can result in routing loops and black-holes that can be very difficult to troubleshoot**

Using OSPF Authentication

- Use authentication; too many people overlook this basic feature
- When using authentication, use the MD5 feature

  area <area-id> authentication message-digest (whole area)

  ip ospf message-digest-key 1 md5 <key>

- Authentication can selectively be disabled per interface with:

  ip ospf authentication null

Agenda

Cisco.com

- Pre-Requisite
- Link State Protocol Refresher
- Fundamental Deployment and Analysis
- OSPF Configuration Examples
- **ISIS Configuration Examples**

How To Configure ISIS?

R1 Configuration

```
!
interface Loopback0
 ip address 172.16.1.1 255.255.255.255
!
interface Ethernet0
 ip address 172.16.12.1 255.255.255.0
 ip router isis
!
router isis
 passive-interface Loopback0
 net 49.0001.1720.1600.1001.00
!
```

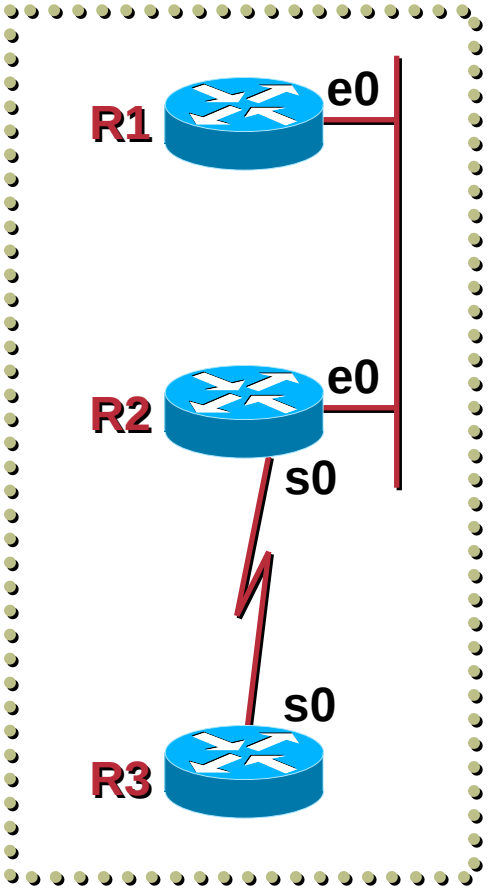

How To Configure ISIS? (Cont.)

R2 Configuration

```
!
interface Loopback0
 ip address 172.16.2.2 255.255.255.255
!
interface Ethernet0
 ip address 172.16.12.2 255.255.255.0
 ip router isis
!
interface Serial0
 ip address 172.16.23.1 255.255.255.252
 ip router isis
!
router isis
 passive-interface Loopback0
 net 49.0001.1720.1600.2002.00
!
```

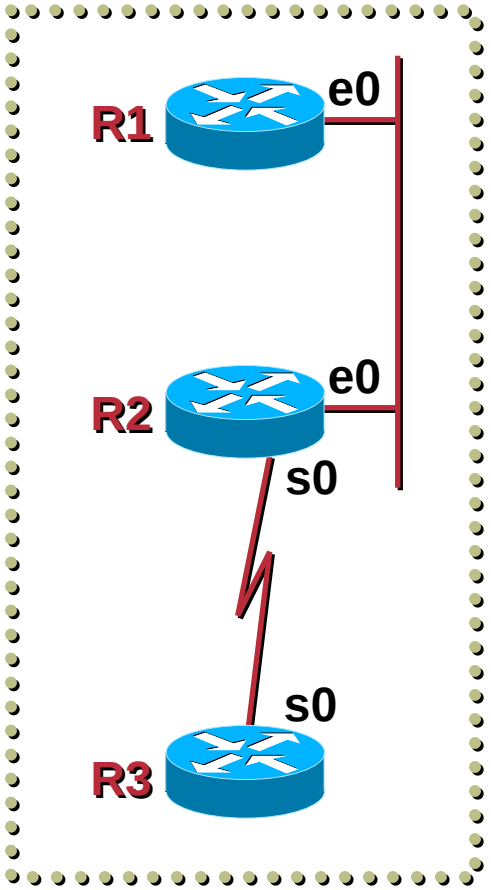

Looking at the Show Commands

Cisco.com

```
R1#show clns neighbor
System Id   Interface  SNPA            State  Holdtime  Type Protocol
R2          Et0        0000.0c47.b947  Up     24        L1L2 IS-IS
```

```
R1#show clns interface ethernet 0
Ethernet0 is up, line protocol is up
  Checksums enabled, MTU 1497, Encapsulation SAP
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
    Interface number 0x0, local circuit ID 0x1
    Level-1 Metric: 10, Priority: 64, Circuit ID: R2.01
    Number of active level-1 adjacencies: 1
    Level-2 Metric: 10, Priority: 64, Circuit ID: R2.01
    Number of active level-2 adjacencies: 1
    Next IS-IS LAN Level-1 Hello in 5 seconds
    Next IS-IS LAN Level-2 Hello in 1 seconds
```

Looking into the Database

Cisco.com

```
R2#show clns neighbor
System Id  Interface  SNPA            State  Holdtime  Type Protocol
R1         Et0        0000.0c09.9fea  Up     24        L1L2 IS-IS
R3         Se0        *HDLC*          Up     28        L1L2 IS-IS
```

```
R2#show isis database
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x0000008B   0x6843        55                0/0/0
R2.00-00            * 0x00000083   0x276E        77                0/0/0
R2.01-00            * 0x00000004   0x34E1        57                0/0/0
R3.00-00              0x00000086   0xF30E        84                0/0/0
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000092   0x34B2        41                0/0/0
R2.00-00            * 0x0000008A   0x7A59        115               0/0/0
R2.01-00            * 0x00000004   0xC3DA        50                0/0/0
R3.00-00              0x0000008F   0x0766        112               0/0/0
```

Looking into the Database Detail

Cisco.com

```
R2#show isis  database R2.00-00 detail
IS-IS Level-1 LSP R2.00-00
LSPID             LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
R2.00-00        * 0x00000093   0x077E        71            0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  Hostname: R2
  IP Address:   172.16.2.2
  Metric: 10         IP 172.16.12.0 255.255.255.0
  Metric: 0          IP 172.16.2.2 255.255.255.255
  Metric: 10         IP 172.16.23.0 255.255.255.252
  Metric: 10         IS R2.01
  Metric: 10         IS R3.00
IS-IS Level-2 LSP R2.00-00
LSPID             LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
R2.00-00        * 0x0000009A   0x5A69        103           0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  Hostname: R2
  IP Address:   172.16.2.2
  Metric: 10         IS R2.01
  Metric: 10         IS R3.00
  Metric: 10         IP 172.16.23.0 255.255.255.252
  Metric: 10         IP 172.16.1.1 255.255.255.255
  Metric: 10         IP 172.16.3.3 255.255.255.255
  Metric: 0          IP 172.16.2.2 255.255.255.255
  Metric: 10         IP 172.16.12.0 255.255.255.0
```

Looking into the Routing-Table

Cisco.com

```
R1#show ip route isis
i L1  172.16.2.2/32 [115/10] via 172.16.12.2, Ethernet0
i L1  172.16.3.3/32 [115/20] via 172.16.12.2, Ethernet0
```

```
R2#show ip route isis
i L1  172.16.1.1/32 [115/10] via 172.16.12.1, Ethernet0
i L1  172.16.3.3/32 [115/10] via 172.16.23.2, Serial0
```

Summary

Cisco.com

- Pre-Requisite
- Link State Protocol Refresher
- Fundamental Deployment and Analysis
- OSPF Configuration Examples
- ISIS Configuration Examples

Recommended Reading

- **Cisco OSPF Command and Configuration Handbook**
  ISBN: 1-58705-071-4
- **Large-Scale IP Network Solutions**
  ISBN: 1-57870-084-1
- **OSPF Network Design Solutions**
  ISBN: 1-57870-046-9

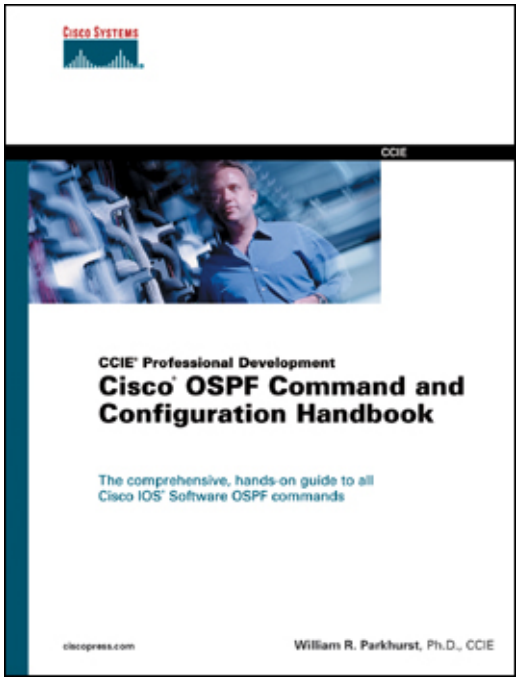

Available On-Site at the Cisco Company Store

Recommended Reading

- **Troubleshooting IP Routing Protocols**
  ISBN: 1-58705-019-6
- **CIM IP Routing: Link-State Protocols**
  ISBN: 1-58720-036-8
- **Routing TCP/IP, Vol. I**
  ISBN: 1-57870-041-8

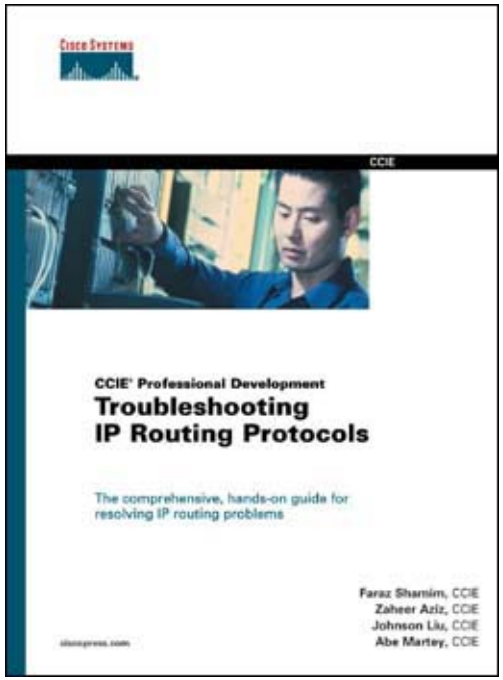

Available On-Site at the Cisco Company Store

NETWORKERS 2003

THE POWER TO TRANSFORM BUSINESS. now.

Please Complete Your Evaluation Form

Session RST-2002