# Intermediate System to Intermediate System

ROUTE Module 9

# Agenda

- OSI Introduction
- IS-IS Basics
- Configuration
- Troubleshooting

# OSI Networks Intro

# OSI Networks

- Interesting computer networks evolution during 70s/80s:
  - US DoD developed Network Control Program that was later replaced by TCP/IP stack
  - Meantime ISO tried to develop and introduce own ISO/OSI stack
- TCP/IP was successful because of easy understanding and quick implementation
- ISO delayed itself by unnecessary formalisms, high-level abstract protocols/interfaces and slow standardization process
  - *ISO protocols are complicated!*
  - Nowadays ISO could be still found in telcos networks but they have never been widely deployed

# OSI Protocols

| OSI Reference Model | OSI Protocol Suite |
|---|---|
| Application | CMIP, DS, FTAM, MHS, VTP<br>ASES: ACSE, ROSE, RTSE, CCRSE, ... |
| Presentation | Presentation Service/Presentation Protocol |
| Session | Session Service/Session Protocol |
| Transport | TP0, TP1, TP2, TP3, TP4 |
| Network | IS-IS, CONP/CMNS, ES-IS, CLNP/CLNS |
| Data Link | IEEE 802.2, IEEE 802.3, IEEE 802.5/ Token Ring, FDDI, X.25 |
| Physical | IEEE 802.3 Hardware, Token Ring Hardware, FDDI Hardware, X.25 Hardware |

# OSI Routing

- Multiple routing protocols were proposed for OSI
  - ES-IS: routing between end-station and its gateway
  - IS-IS: routing between routers in one AS a.k.a. domain in ISO terminology
  - IDRP: routing between domains (analogous to BGP)
- Properties of IS-IS turn out to be sophisticated and flexible
  - IS-IS was proposed and functional before OSPF, OSPF started as just a lite version of IS-IS
  - During migration from OSI to IP it was suitable to have routing protocol capable of using both stacks
  - RFC 1195 integrated extension to cooperate with IP without redefining basic structure of protocol

# OSI Terms

- **End System (ES)**: end station, host (PC)
- **Intermediate System (IS)**: router
- **Area**: the set od interconnected ESs and ISs sharing same topology information
- **Domain**: the set of interconnected areas (same as AS)
- Routing levels:
  - Level 0: routing among ES and IS
  - Level 1: routing inside the one area
  - Level 2: routing among areas inside the one domain
  - Level 3: routing among domains

# Routing Levels

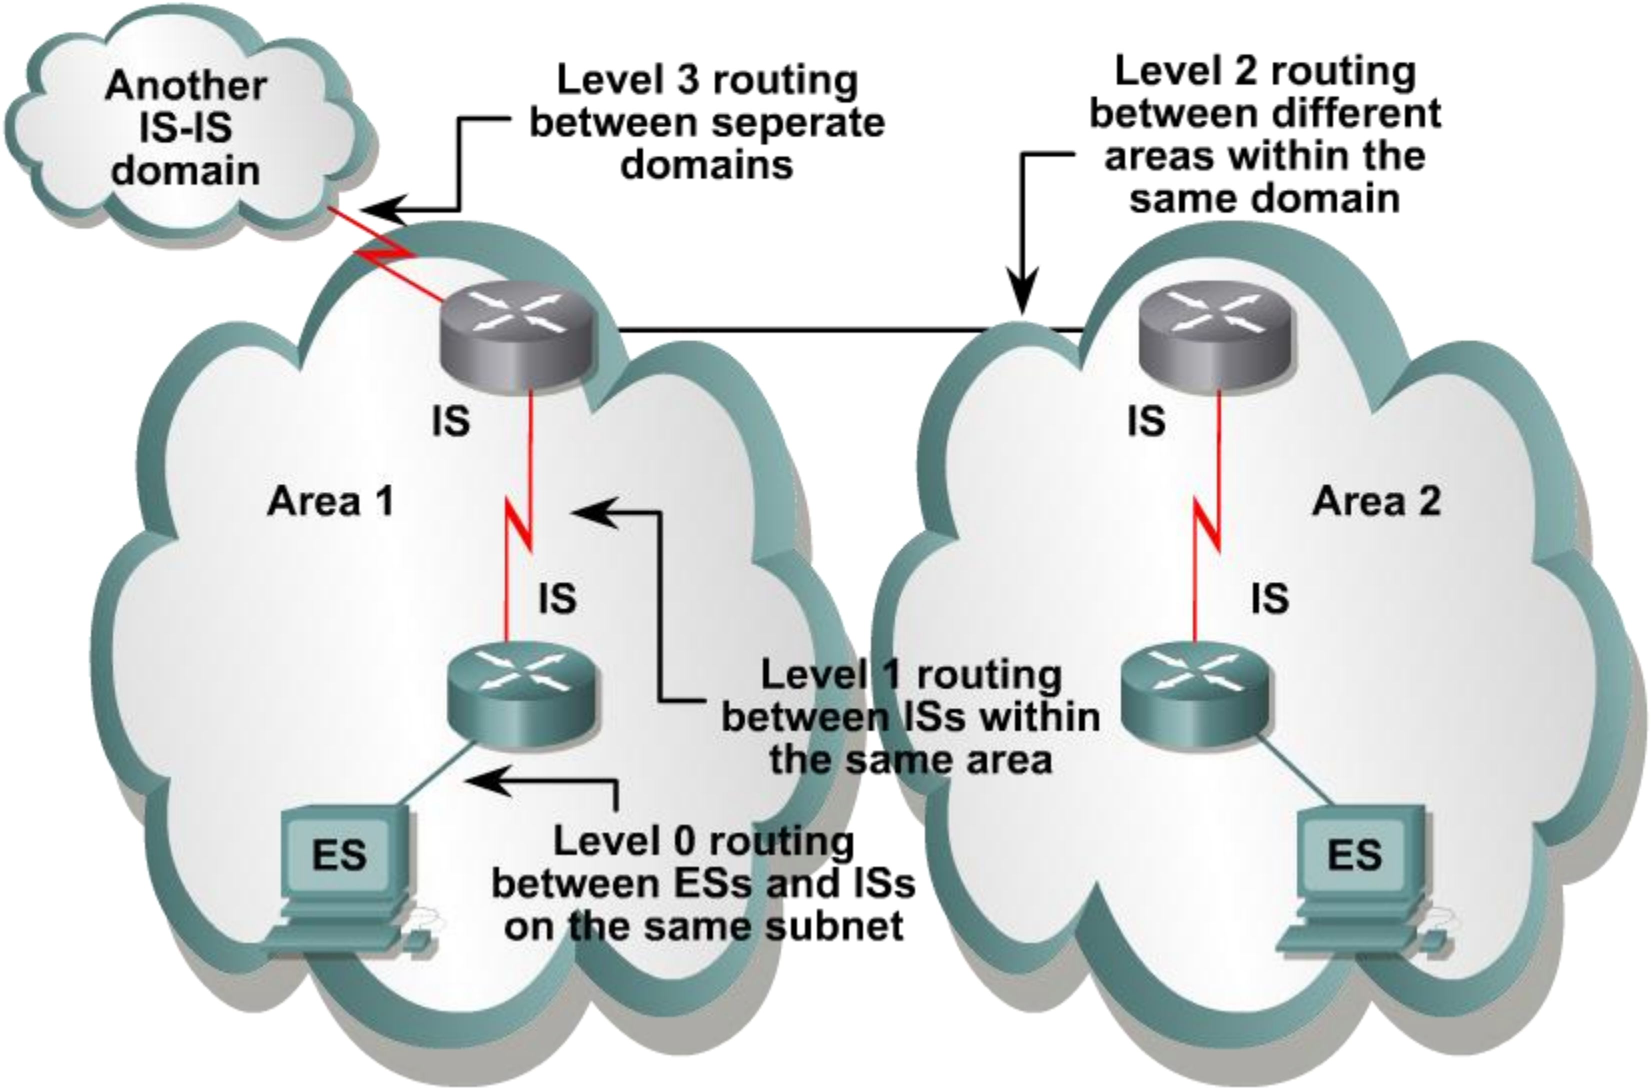

# Node Identifiers and Its Interfaces

- Node address in OSI networks contains domain number, area number, node identifier and particular service on it
  - Node has L3 address as a unit not as IP address for each interface
- **Network Service Access Point (NSAP)**
  - OSI address name
  - Abstract point between network and transport layers
  - NSAP has variable length in range from 8 to 20 B
- **NSAP Selector (NSEL)**
  - The last byte of NSAP
  - Identifier of network service
  - `IF` NSEL == 0 `THEN` it is called **Network Entity Title (NET)**
  - Only NET are assigned when configuring IS-IS router

# NSAP Address

The Cisco implementation supports all NSAPs defined by ISO 8348/Addendum 2 and ISO 10589

**NSAP Addressing Structure**

| IDP | | DSP | | |
|---|---|---|---|---|
| AFI | IDI | AREA | STATION ID | SEL |

- The NSAP is formed by the Inter-Domain Part (IDP) and the Domain Specific Part (DSP)
- The IDP is composed of:
  - One-byte Authority and Format Identifier (AFI)
  - Variable Length Initial Domain Identifier (IDI)
- The DSP is composed of:
  - Two-byte area
  - Six-byte station ID
  - One-byte selector

# Example: NET Assignment

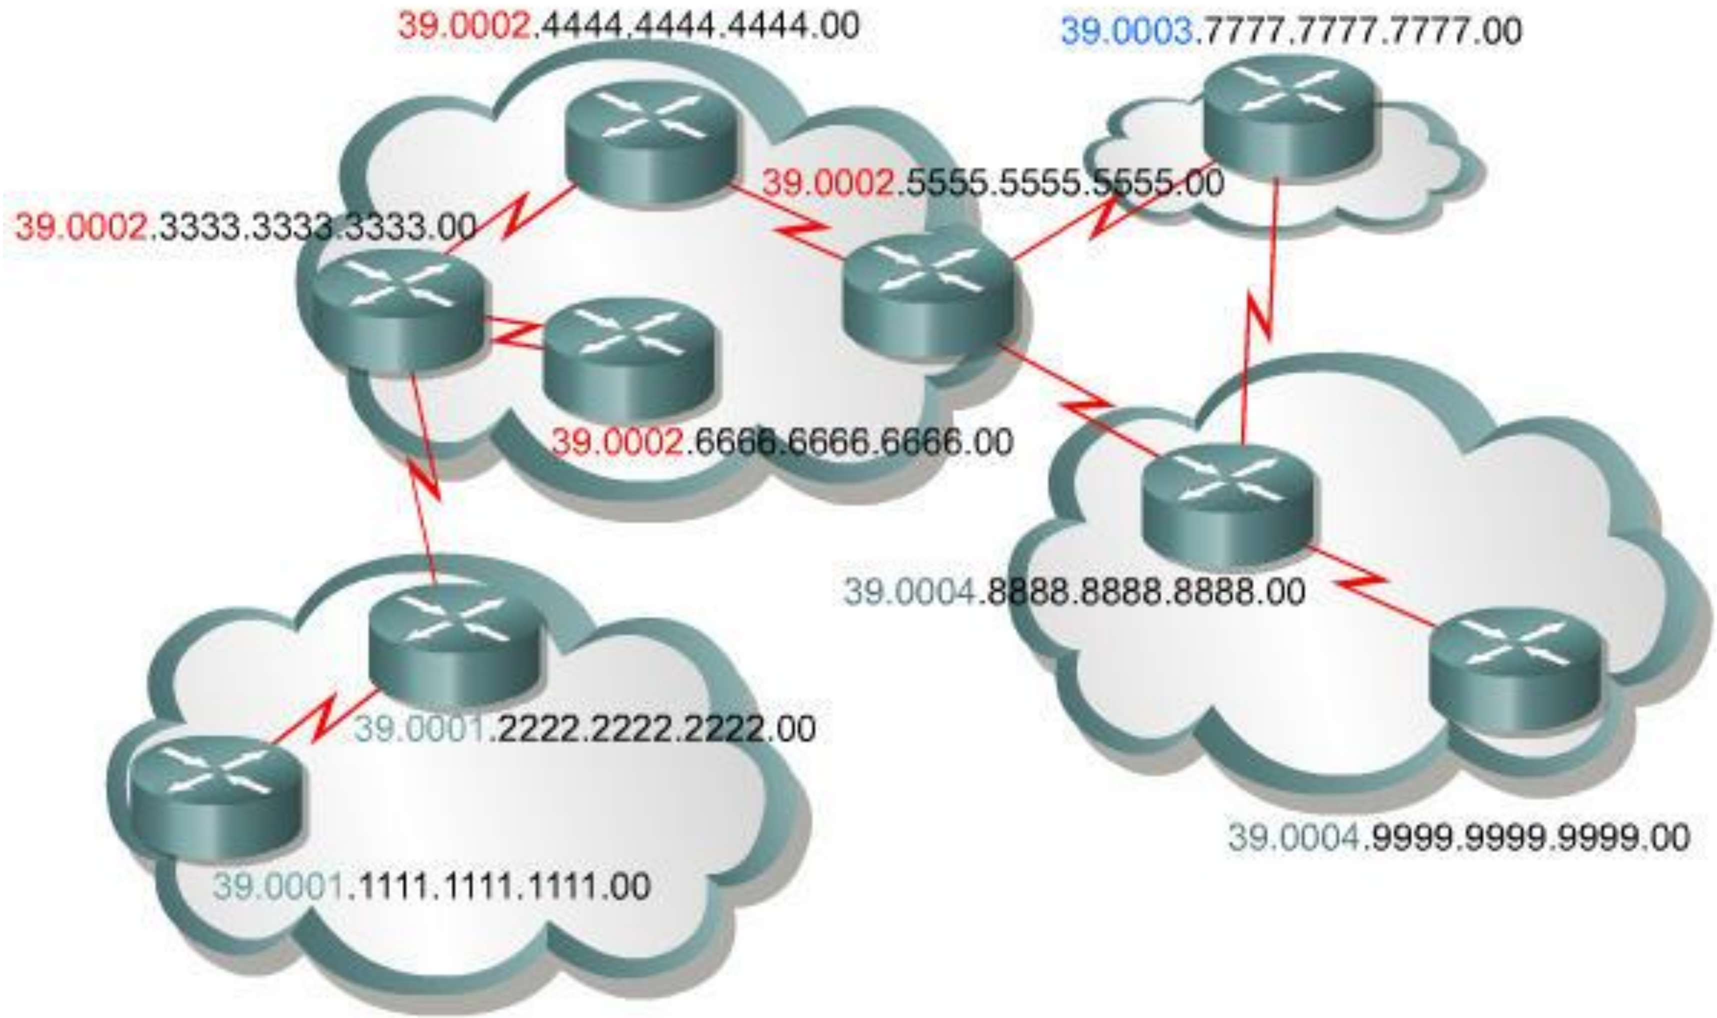

# Node Address

- NSAP/NET addresses should be read from right to left
  **49.0001.1122.3344.5566.00**
  - The most right byte: **NSEL**
  - The next 6 bytes: **System ID**
  - The remaining bytes except the last one: **HO-DSP, IDI** – their length and semantics is specified by AFI
  - The most left byte: **AFI**, for private domains reserved 49
- System ID MUST BE unique…
  - …among Level1 routers inside one area
  - …among Level2 routers inside one domain
  - *The best case is when System ID is unique for any router inside domain*

# Node Interface

- Each and every node interface is identified by **SubNetwork Point of Attachment (SNPA)**
  - L2 identity of interface
  - Ethernet: MAC address
  - Frame Relay, ATM, X.25: DLCI
  - HDLC and PPP
- Router mark each interface with **Circuit ID** for internal purposes
  - 1 B long number (some IS-IS extensions uses larger space)
  - Assigned by system itself automatically and it CAN NOT be changed via any configuration
  - Multiaccess segment has number that is composed of System ID of Designated IS (analogy of OSPF DR) and its Circuit ID for this segment

# IS-IS Basics

# IS-IS Routing Protocol

- IS-IS is link-state protocol just like OSPF, except it is completely different (and better) than OSPF ☺
- IS-IS was originally designed for OSI networks
- Later it was integrated with IPv4 and IPv6 support
- Current IS-IS implementation supporting multiple address families are called Integrated IS-IS
- Integrated IS-IS is classless for IPv4 and it supports summarization, authentication, fast-convergence, area partitioning and much more

# Metrics

- IS-IS defines 4 different types of metric
  - Default
  - Expense (financial costs for data transfer across the link)
  - Delay
  - Error (error rate on the link)
- Cisco implementation supports only default metric type
  - Moreover, default value of this Cisco implementation Default metric is not connected to any property of the link
  - On every interface the default value is 10
  - Hence, it is wise to change it according to real situation

# Messages

- IS-IS messages are carried directly in data-link frames
  - It does not use any transport protocol!
  - It is address family agnostic which means that it is independent on transport protocol

| IS-IS Adjacencies | |
|---|---|
| AllL1ISs | 01-80-C2-00-00-14 |
| AllL2ISs | 01-80-C2-00-00-15 |
| AllIntermediateSystems | 09-00-2B-00-00-05 |
| AllEndSystems | 09-00-2B-00-00-04 |

# Message Types

- IS-IS has four message types
  - **Hello packet (IIH)**
  - **Link-state PDU (LSP)**
  - **Complete Sequence Number PDU (CSNP)**
  - **Partial Sequence Number PDU (PSNP)**
- **Hello packet** a.k.a **IIH (IS-IS Hello)**
  - Sent every 10 seconds
  - Designated IS sends it 3× faster (by default every 3.3 seconds)
  - Timeout is 3× Hello interval
  - There are two different Hello packets for Level1 and Level2 routing multiaccess networks
  - Only one type of Hello paket for point-to-point networks
  - Timers DO NOT need to match across area
- In RFC are mentioned also ESH and ISH – they are exchanged between ES and IS

# Hello Packets

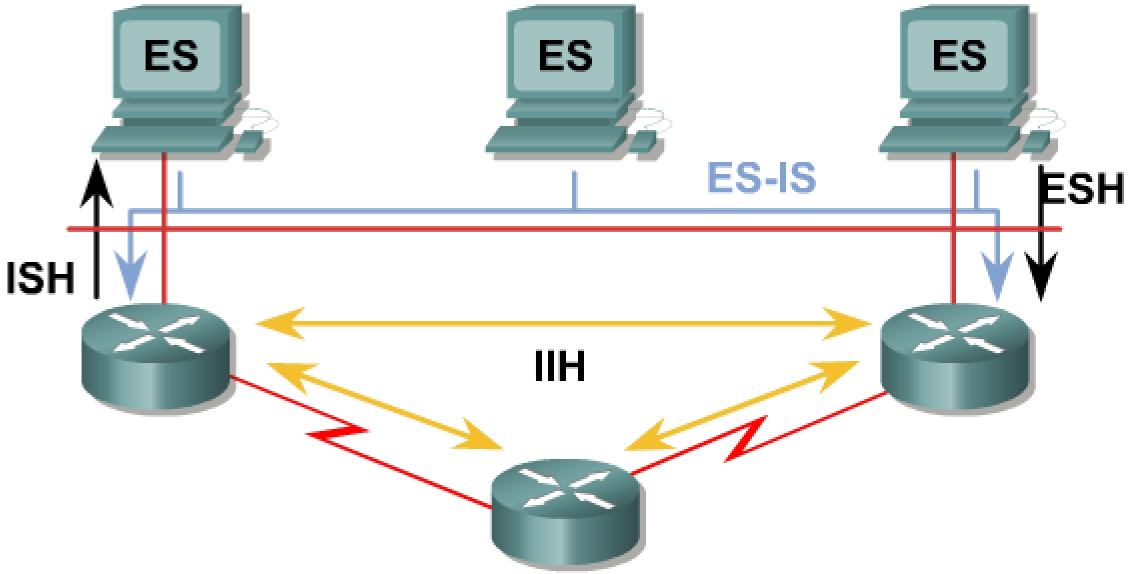


- End-stations inform routers about their existence with ESH (ES Hello) in OSI networks
  - *Same thing DOES NOT exist in IP networks!*
- Routers send ISH (IS Hello) to inform end-station about their existence
- Routers exchange IIH (IS-IS Hello)

# LSP

- **Link-State PDU / Link-State Packet (LSP)**
  - Generated by each router (IS), similar to LSA from OSPF
  - Contains topological information about router, its interfaces, ES available through this router. For IP networks also list of directly connected IP networks
  - LSPs are numbered from 0x00000001 to 0xFFFFFFFF (4B)
  - LSP´s Lifetime is 1200s (20min) – `max-lsp-lifetime` - and it gets periodically refreshed every 900s (15min) `lsp-refresh-interval`
  - In the event of LSP aging out, it is kept in database for 1 minute (ZeroAgeLifetime)
  - After LSP´s sequence number reaches upper limit (0xFFFFFFFF) its originator must wait 21minutes for the LSP to expire and get remove from database. Then it can continue with sequence number 1.

# PSNP and CSNP

- **Partial/Complete Sequence Number PDUs**
  - Similar function to DDP/LSR/LSAck packets from OSPF
  - On broadcast networks:
    - “implicit acknowledgments”
    - DIS sends CSNP every 10 seconds
    - PSNP is used for requesting newer LSP
  - On point-to-point networks:
    - CSNP is send after adjacency state change to UP
    - IS sends LSPs that are missing in neighbors LSP-DB
    - PSNP acknowledge received LSP

# IS-IS Network Types

- IS-IS natively supports following network types
  - Broadcast e.g. Ethernet
  - Point-to-Point e.g. PPP
- IS-IS does not have special support for NBMA
- NBMA network could be transformed to multiple point-to-point connections
  - E.g. Create special subinterface with own IP network for each DLCI as in case of Frame Relay
- IS-IS could be run over NBMA but full mesh connectivity MUST be guaranteed

# IS-IS and broadcast networks

- Similar to OSPF, IS-IS elects Designated IS (DIS) on broadcast networks
  - Router with highest IS-IS priority is elected DIS.
  - Default priority is 64, but can be changed to value from range 0 – 127
  - In the event of tie, IS with highest SNPA gets elected DIS
    - Not all interfaces have SNPA, or we do not know neighbor's SNPA
  - In the event of tie, IS with the highest SystemID is elected DIS
  - Election is preemptive and happens every time when IS with higher priority appears.
  - There is no backup DIS
- DIS's responsibilities are
  - Create network's LAN ID (SystemID.CircuitID)
  - Generate CSNPs
  - Create pseudonode
  - Unlike OSPF, all ISs on segment are "friends" (adjacent), synchronization is done via CSNP

# Pseudonode for Broadcast Networks ①

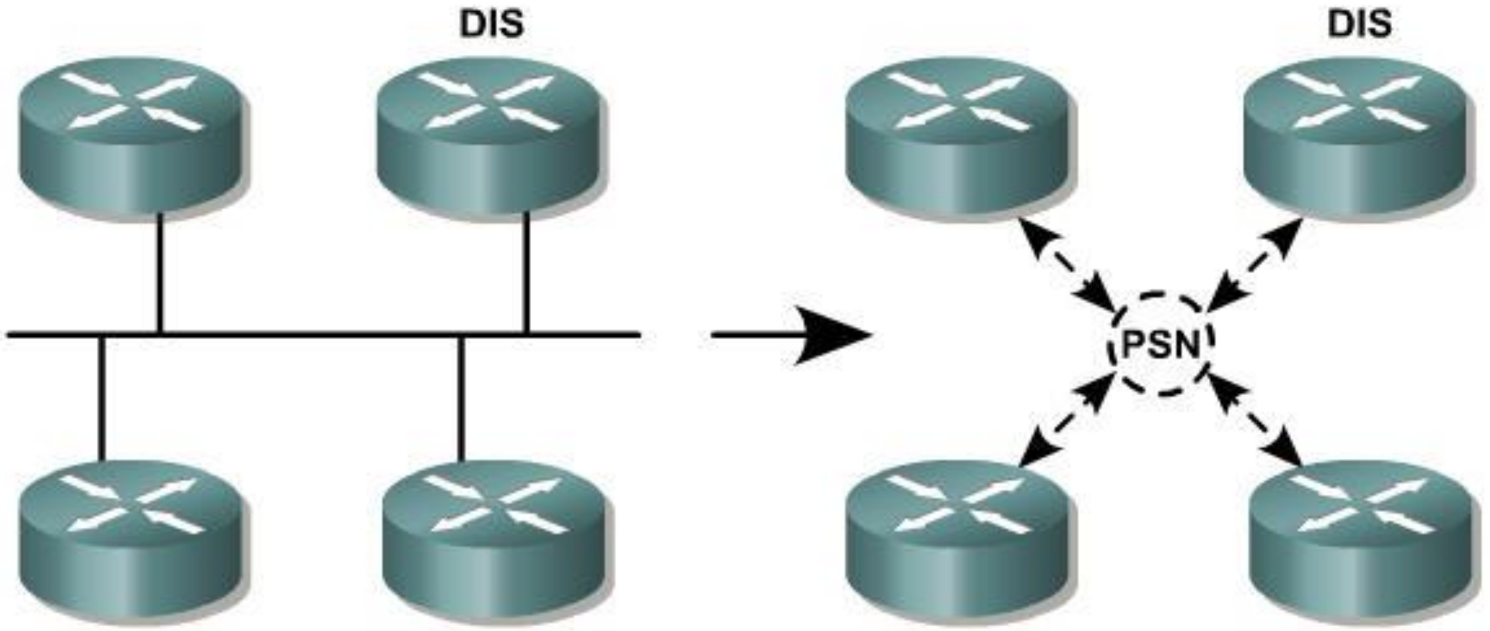

# Pseudonode for Broadcast Networks ②

- Pseudonode is fictive network node that abstracts topology on broadcast segment
  - All $n$ neighbors on segment communicates with this pseudonode ($n$ separate connections) instead of creating $\frac{n.n-1}{2}$ connections with each other
  - DIS generates LSP describing pseudonode apart from its own LSPs
    - This approach is used also by OSPF
    - Pseudonode LSP is similar to OSPF LSA2

# Level of routing in IS-IS

- IS-IS provides routing on Level1 and Level2
  - IS sends also ISH, but they are not responsibility of IS-IS protocol
- There are separate PDUs and LSP-DB for each level
- Level1
  - Routing inside area
  - Contains topological information about area itself – creates graph of nodes using SystemID and locates shortest path within the graph
  - For IP: networks in current area
- Level2
  - Routing between areas
  - Exchanges area addresses (prefixes) and locates shortest paths to them
  - For IP: networks and summary networks from all areas

# Level of routing in IS-IS

- Router can be Level1, Level2 or Level1-2
  - According DBs and adjacencies are created
  - Adjacency is established between neighbors only on matching levels
  - Neighboring routers have to share at least one level to successfully establish adjacency
- Level2 routers must form continuous ‘chrbticu‘
  - Serves as inter-area routing
  - ‘chrbtica’ can be only one

# Areas in IS-IS

- Area identification is part of the NSAP/NET
- Since NET is configured per Router and not per interface, the router as a whole belongs to specified area
  - Therefore the area boundaries are on links
  - Router can be part of up to 3 areas (Cisco routers allows max 3 NETs)
  - It is used for migration between CLNS address spaces
  - Databases for each NET are merged together
- Areas must be connected with continuous chain of Level2 routers

# Areas in IS-IS

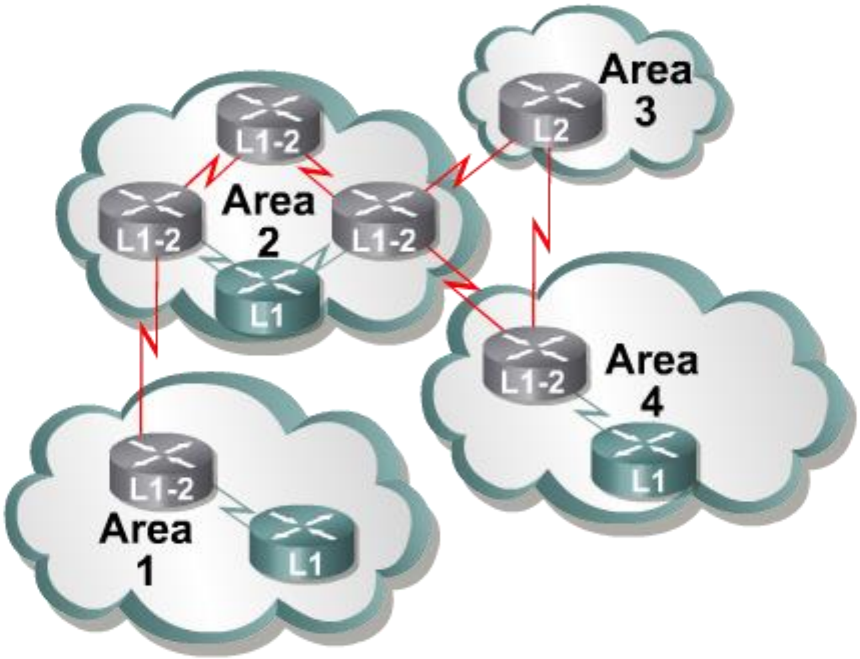


- Level1 routers exchange topology information about its own area
- Level2 routers in its LSPs announce areas they know and connect
- Level1-2 routers serve as border routers between Level1 and Level2 routing

# Areas in IS-IS

- IS-IS Areas do not have types as in OSPF, but their behavior resemble some
  - Level1 routers in areas have information only about its own area. Networks from any other area will be replaced by default route. Level1 can import other directions. This behavior resembles NSSA Totally Stubby
  - Level1-2 routers have information about its own area as well as about networks in other areas. This behavior resembles ABR routers.
  - Level2 routers know networks from all areas. This resembles regular area.
- Carrying IP routing information
  - Level2-capable router adds to its Level2 LSP-DB all directly connected IP networks and all IP networks from its own area

# IS-IS vs. OSPF

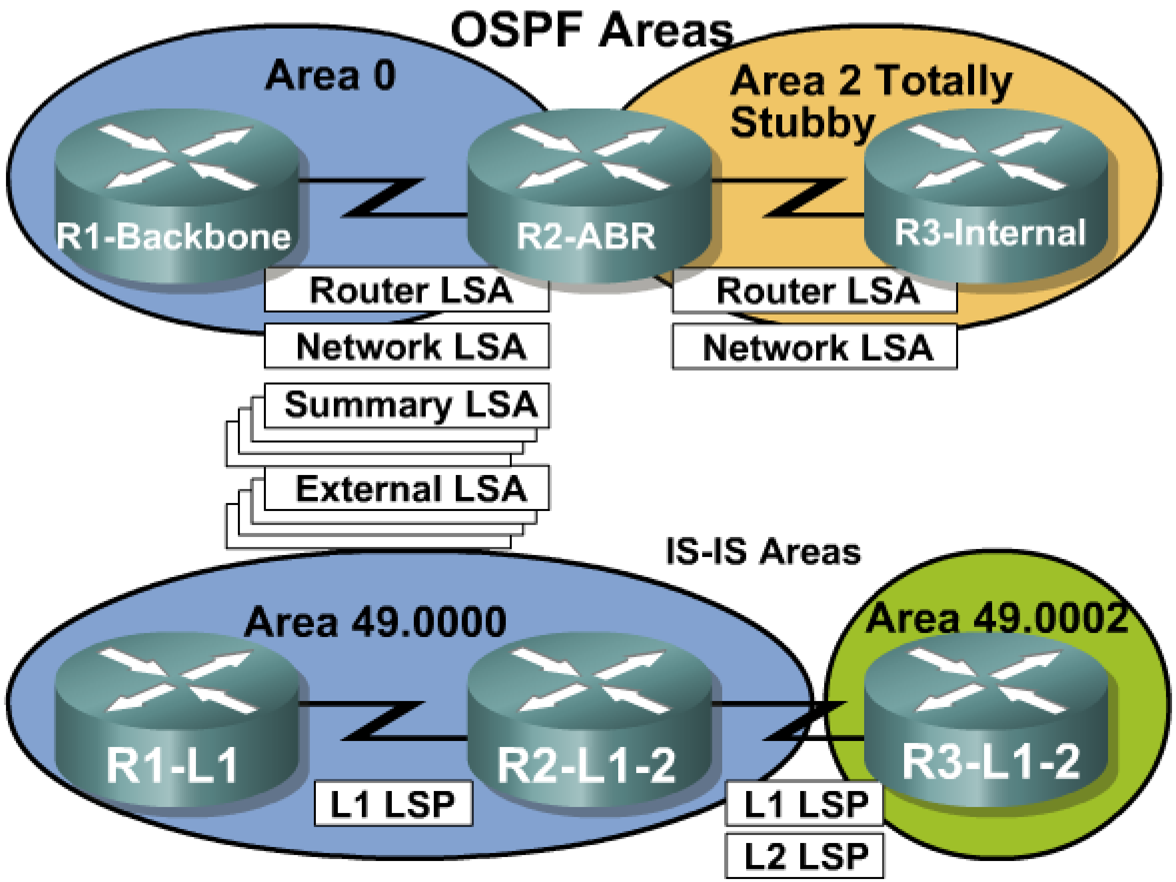
OSPF Areas
Area 0
Area 2 Totally Stubby
R1-Backbone
R2-ABR
R3-Internal
Router LSA
Network LSA
Summary LSA
External LSA
Router LSA
Network LSA
IS-IS Areas
Area 49.0000
Area 49.0002
R1-L1
R2-L1-2
R3-L1-2
L1 LSP
L1 LSP
L2 LSP

# Comparison of OSPF and IS-IS

| OSPF | Integrated IS-IS |
|---|---|
| Area borders are on routers (ABR) | Area borders are on links |
| Each link is exactly in 1 area | Each router is exactly in one area |
| Difficult to extend ´Chrbticu´ | Easy to extend ´Chrbticu´ |
| Produces many small LSA | Produces fewer LSPs |
| Runs on top of IP | Runs directly on Link layer |
| Requires IP addresses | Requires IP and CLNS addresses |
| Metric is proportional to transmission speed of interface | Standard metric is 10 on all interfaces |
| Difficult to extend | Easy to extend through the Type, Length, Value (TLV) mechanism |
| Devices, personal and information are available | Devices, personal and information are less available |

# Configuration

# Requires NET Addresses

- A Common CLNS parameters (NET) and area planning are still required even in an IP environment
- Even when Integrated IS-IS is used for IP routing only, routers still establish CLNS adjacencies and use CLNS packets

# OSI Area Routing: Building an OSI Forwarding Database (Routing Table)

- When databases are synchronized, Dijkstra’s algorithm (SPF) is run on the LSDB to calculate the SPF tree
- The shortest path to the destination is the lowest total sum of metrics
- Separate route calculations are made for Level 1 and Level 2 routes in Level 1-2 routers
- Best paths are placed in the OSI forwarding database (CLNS routing table)

# Building an IP Routing Table

- Partial Route Calculation (PRC) is run to calculate reachability
  - Since IP and ES are represented as leaf objects, they do not participate in SPF
- Best paths are placed in the IP routing table following IP preferential rules
  - They appear as Level 1 or Level 2 IP routes

# Integrated IS-IS Configuration Steps

1. Define areas, prepare addressing plan (NETs) for routers, and determine interfaces
2. Enable IS-IS on the router
3. Configure the NET
4. Enable Integrated IS-IS on the appropriate interfaces. Do not forget interfaces to stub IP networks, such as loopback interfaces (although there are no CLNS neighbors there)

- These are each explained in the next few slides.

# Integrating IS-IS Routing

| Command | Description |
|---|---|
| **router isis** *[area-tag]* | Enables IS-IS as an IP routing protocol and assigns a tag to the process (optional). Given in the global configuration mode. |
| **net** *network-entity-title* | Identifies the router for IS-IS by assigning a NET to the router. Given in the router configuration mode. |
| **clns router isis** [area-tag] | Specifies that the interface is actively routing IS-IS when the network protocol is ISO-CLNS, and identifies the area associated with this routing process on this interface. |
| **ip router isis** *[area-tag]* | Enables IS-IS on the interfaces that run IS-IS. (This approach is slightly different from most other IP routing protocols, where the interfaces are defined by network statements; there is no network statement under the IS-IS process.) Given in interface configuration mode. |
| **ipv6 router isis** *[area-tag]* | Enables the specified IPv6 IS-IS routing process on an interface. |

# Step 1: Define Area and Addressing

- Area determined by NET prefix:
  - Assign to support two-level hierarchy
- Addressing:
  - IP: Plan to support summarization.
  - CLNS: Prefix denotes area. System ID must be unique
- Recommended way of generating SystemID:
  - Create Loopback address on each router
  - Convert IP address of Loopback:

10.15.134.7 → 010.015.134.007 → 0100.1513.4007

# Step 2: Enable IS-IS on the Router
# Step 3: Configure the NET

```
Router(config)#
router isis [area-tag]
```

- Enable the IS-IS routing protocol.

  ***area-tag*** – name for a process

- When routing of CLNS packets is also needed, use the **clns routing** command.

```
Router(config-router)#
net network-entity-title
```

- Configure an IS-IS NET address for the routing process.

# Step 4: Enable Integrated IS-IS

```
router(config-if)# ip router isis [area-tag]
router(config-if)# ipv6 router isis [area-tag]
```

- Includes an interface in an IS-IS routing process
- There is no **network** command for IS-IS!!!
- Configuration snippet:

```
interface FastEthernet0/0
 ip address 10.1.1.2 255.255.255.0
 ip router isis
!
interface Serial 0/0/1
 ip address 10.2.2.2 255.255.255.0
 ip router isis
!
router isis
 net 49.0001.0000.0000.0002.00
```

# Simple Integrated IS-IS Example

The configured router acts as an IP-only Level 1-2 router.

```
interface FastEthernet0/0
 ip address 10.1.1.2 255.255.255.0
 ip router isis
!
interface Serial 0/0/1
 ip address 10.2.2.2 255.255.255.0
 ip router isis
!

<output omitted>

router isis
 net 49.0001.0000.0000.0002.00
```

# Change IS-IS Level

```
Router(config-router)#
is-type {level-1 | level-1-2 | level-2-only}
```

- Configure the IS-IS level globally on a router; the default is level 1-2.

```
Router(config-if)#
isis circuit-type {level-1 | level-1-2 | level-2-only}
```

- Configure the type of adjacency on an interface; the default is Level 1-2.

# Change IS-IS Metric

```
Router(config-if)#
isis metric metric [delay-metric [expense-metric [error-metric]]]
{level-1 | level-2}
```

- Configure the metric for an interface; the default is 10
- Metric value is from 1 to 63

```
Router(config-router)#
metric default-value {level-1 | level-2}
```

- Alternately, configure the metric globally for all interfaces

# Narrow and Wide Metrics

- Former IS-IS specification uses metrics that are:
  - for interfaces 6 bits long (values in range from 1 to 63, 0 is reserved)
  - for a whole path 10 bits long (max. value 1023)
- Those kind of so called **narrow metrics** are unsuitable for nowadays networks
- Currently IS-IS supports also **wide metrics** that are:
  - 24 bits long for interface
  - 32 bits long for whole path (some bits are reserved)
- It is suggested to enable wide metrics, but they should be configured on all IS-IS devices in the domain
  - *Otherwise routing loops might occur!!!*

# Changing Metric Types

```
Router(config-router)#
metric-style {narrow | transition | wide}
   [level-1 | level-2 | level-1-2]
```

- Previous command specifies the IS-IS metric type
- Option **`transition`** is used when migrating from narrow to wide metrics
- All routers in domain MUST use same type of metric

# Tuning IS-IS Configuration

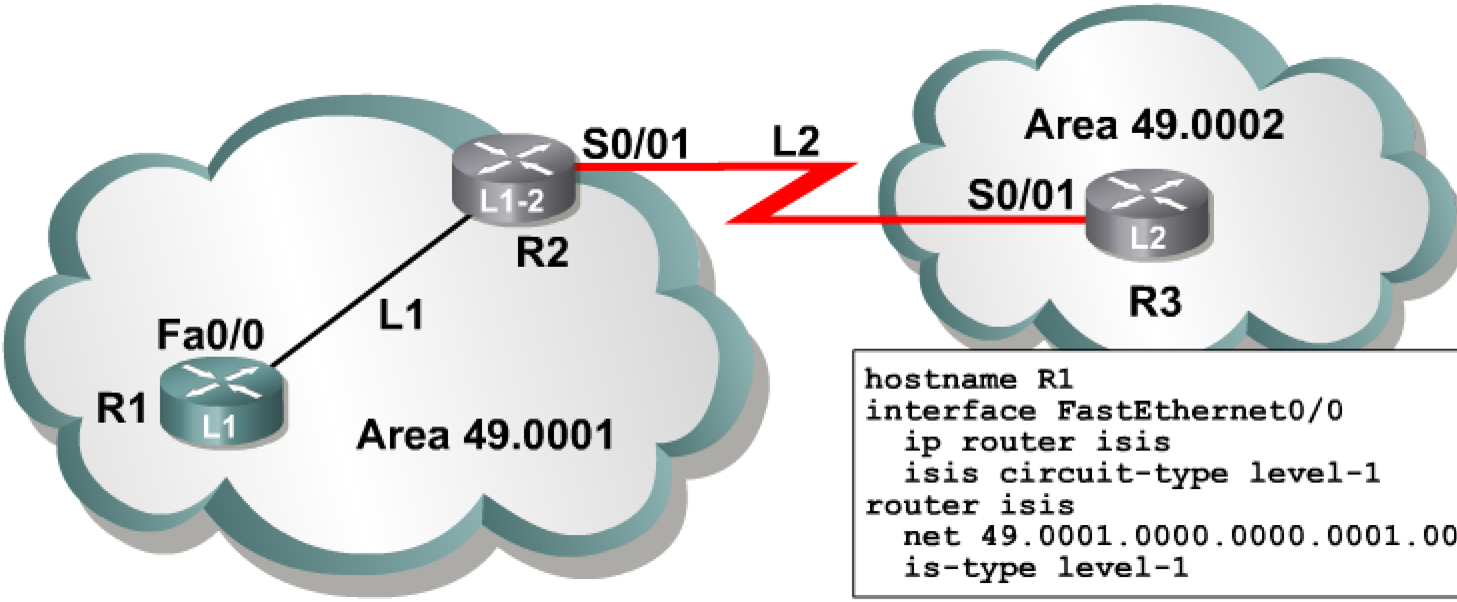


```
hostname R1
interface FastEthernet0/0
  ip router isis
  isis circuit-type level-1
router isis
  net 49.0001.0000.0000.0001.00
  is-type level-1
```

```
hostname R2
interface FastEthernet0/0
  ip router isis
  isis circuit-type level-1
interface Serial0/0/1
  ip router isis
  isis circuit-type level-2-only
  isis metric 35 level-2
router isis
  net 49.0001.0000.0000.0002.00
```

- Change router type on R1 and R3
- Change interface levels on R2
- Change metric on S0/0/1

# IP Summarization

```
Router(config-router)#
summary-address address mask [level-1 | level-2
    |level-1-2][tag tag-number] [metric metric-value]
```

- Creates summary and appropriate discard route
- Command applies only on the router that is actually importing routes into any ISIS topology database:
  - any L1L2 router
  - any router performing redistribution into IS-IS
- Level parameter specifies in which level summary is created
  - Default is Level 2
  - Summary into L1 makes sense only in case of redis
- Configuration snippet to summarize 10.3.2.0/23 into L1-2:

```
P3R1(config-router)#
summary-address 10.3.2.0 255.255.254.0 level-1-2
```

# Authentication: The Older Way

```
interface FastEthernet0/0
  isis password HESLO1 [level-1 | level-2]

router isis
 area-password HESLO2
 domain-password HESLO3
```

- Beware only plain-text password!!!
- The command **isis password** appends password to all Hello packets (authenticates adjacencies)
  - For point-to-point networks is password valid for both L1 and L2
  - For broadcast networks password could be different for L1 and L2
- The command **area-password** defines password authenticating to Level1 LSP, PSNP and CSNP
- The command **domain-password** defines password authenticating Level2 LSP, PSNP and CSNP

# Authentication: The New Way

```
interface FastEthernet0/0
  isis authentication mode {md5 | text} [level-1 | level-2]
  isis authentication key-chain KEYCHAIN

router isis
  authentication mode {md5 | text} [level-1 | level-2]
  authentication key-chain KEYCHAIN
```

- Same way as in case of authentication process for distance vector protocols using keychains
- Per-interface commands specify authentication for Hello packets
- ISIS sub-configuration commands specify authentication for LSP, CSNP, PSNP
- Key numbers and key-string MUST match between neighbors

# Threat of Suboptimal Routing

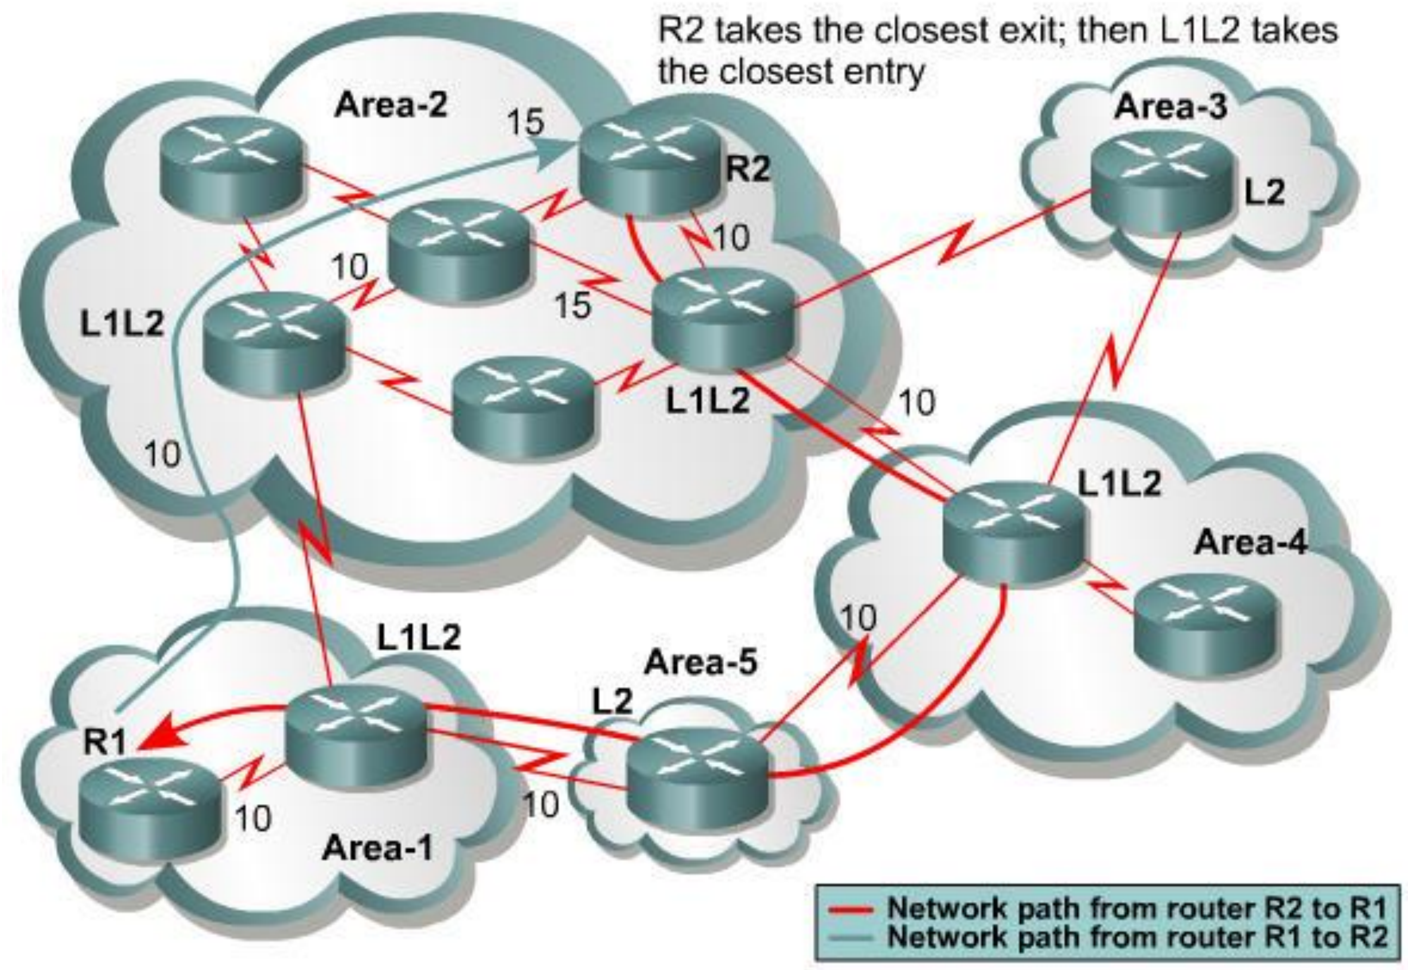

# Suboptimal Routing in IS-IS

- Suboptimal routing happens because of separate LSDB for Level1 and Level2
  - Routers in the one area populate Level1 LSDB
  - Information are transferred from Level1 into Level2 only on border routers residing in both levels
  - Level1 router routes through the closest Level1-2 router
- Solution: **route leaking** = redistribution from L2 to L1
  - `redistribute` command in IS-IS configuration subsection
  - Wide metric is highly recommended when benefiting from route leaking

# Troubleshooting

# Example: Topology

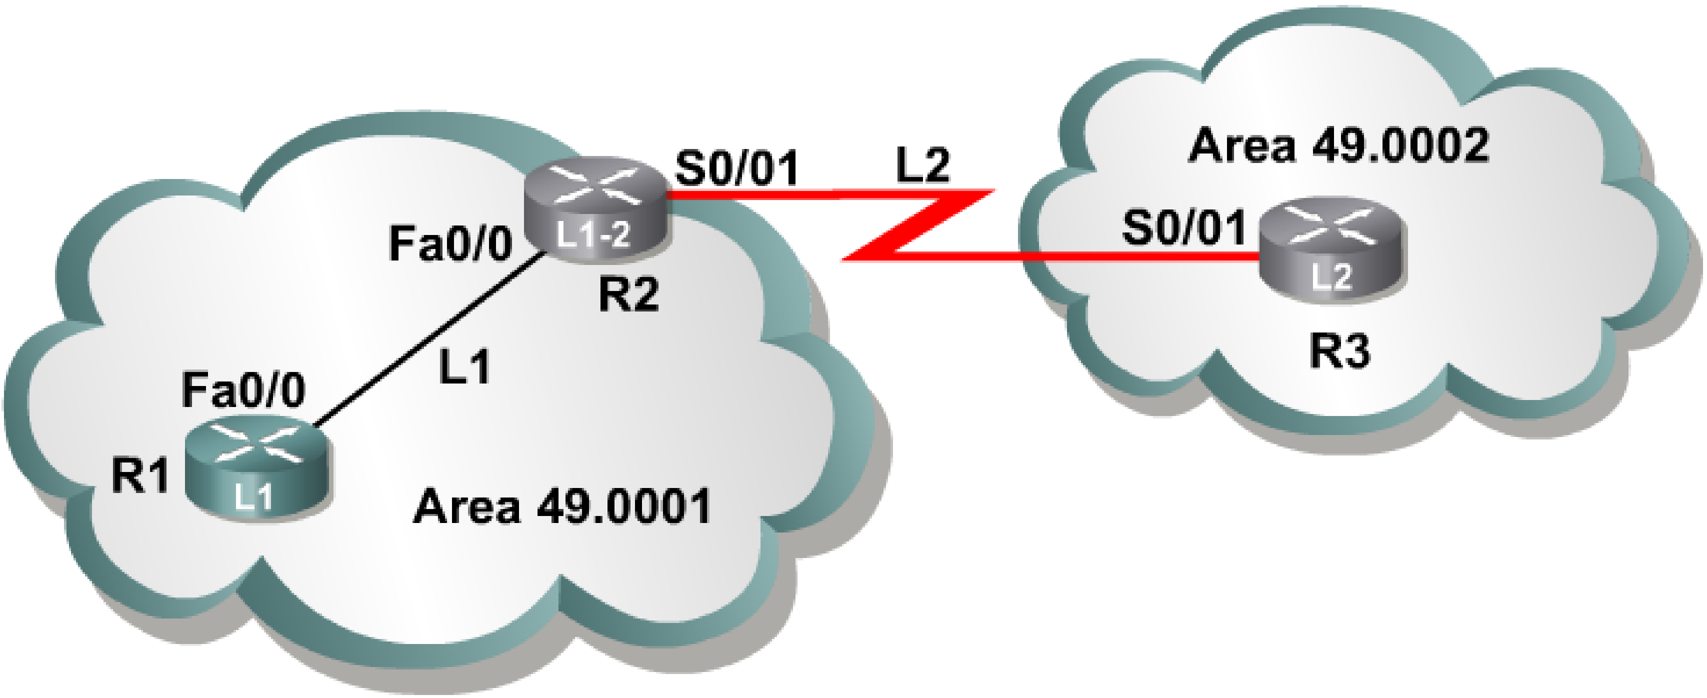

# Example: Is Integrated IS-IS Running?

```
R2# show ip protocols
Routing Protocol is "isis"
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: isis
  Address Summarization:
    None
  Maximum path: 4
  Routing for Networks:
    FastEthernet0/0
    Loopback0
    Serial0/0/1
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.10.10.10          115      00:00:02
    10.30.30.30          115      00:00:03
  Distance: (default is 115)
```

- Displays the parameters and current state of the active routing protocol processes

# Example: Are There Any IP Routes?

```
Router#
show ip route [address [mask]] | [protocol [process-id]]
```

```
R2# show ip route isis
     10.0.0.0/24 is subnetted, 5 subnets
i L2    10.30.30.0 [115/45] via 10.2.2.3, Serial0/0/1
i L1    10.10.10.0 [115/20] via 10.1.1.1, FastEthernet0/0
R2#
```

- Displays the current state of the routing table

# Troubleshooting Commands: CLNS

```
Router#
show clns
```

- Displays information about the CLNS network

```
Router#
show clns [area-tag] protocol
```

- Lists the protocol-specific information

```
Router#
show clns interface [type number]
```

- Lists the CLNS-specific information about each interface

```
Router#
show clns [area-tag] neighbors [type number] [detail]
```

- Displays both ES and IS neighbors

# Troubleshooting CLNS and IS-IS

```
Router#
show isis [area-tag] route
```

- Displays IS-IS Level 1 routing table (system IDs)

```
Router#
show clns route [nsap]
```

- Displays IS-IS routing table (areas)

```
Router#
show isis [area-tag] database
```

- Displays the IS-IS LSDB

```
Router#
show isis [area-tag] topology
```

- Displays IS-IS least-cost paths to destinations

# Example: L1 and L2 Topology Table

```
R1# show isis topology
IS-IS paths to level-1 routers
System Id             Metric     Next-Hop  Interface  SNPA
R1                    --
R2                    10         R2        Fa0/0
0016.4650.c470

R2# show isis topology
IS-IS paths to level-1 routers
System Id             Metric     Next-Hop  Interface  SNPA
R1                    10         R1        Fa0/0
0016.4610.fdb0
R2                    --
IS-IS paths to level-2 routers
System Id             Metric     Next-Hop  Interface  SNPA
R1                    **
R2                    --
R3                    35         R3        Se0/0/1    *HDLC*
```

# Example: What About CLNS Protocol?

```
R2# show clns protocol

IS-IS Router: <Null Tag>
  System Id: 0000.0000.0002.00  IS-Type: level-1-2
  Manual area address(es):
        49.0001
  Routing for area address(es):
        49.0001
  Interfaces supported by IS-IS:
        Loopback0 - IP
        Serial0/0/1 - IP
        FastEthernet0/0 - IP
  Redistribute:
    static (on by default)
  Distance for L2 CLNS routes: 110
  RRR level: none
  Generate narrow metrics: level-1-2
  Accept narrow metrics:   level-1-2
  Generate wide metrics:   none
  Accept wide metrics:     none
```

# Example: Adjacencies

```
R2# show clns neighbors
System Id      Interface   SNPA              State  Holdtime   Type Protocol
R3             Se0/0/1     *HDLC*            Up     28         L2   IS-IS
R1             Fa0/0       0016.4610.fdb0    Up     23         L1   IS-IS

R2#show clns interface s0/0/1
Serial0/0/1 is up, line protocol is up
  Checksums enabled, MTU 1500, Encapsulation HDLC
  ERPDUs enabled, min. interval 10 msec.
  CLNS fast switching enabled
  CLNS SSE switching disabled
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 45 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-2
    Interface number 0x1, local circuit ID 0x100
    Neighbor System-ID: R3
    Level-2 Metric: 35, Priority: 64, Circuit ID: R2.00
    Level-2 IPv6 Metric: 10
    Number of active level-2 adjacencies: 1
    Next IS-IS Hello in 5 seconds
    if state UP
```

# IS-IS Routing Protocol - Restrictions

IS-IS for IPv6 uses the same SPF for both IPv4 and IPv6.

Therefore:

- If using IS-IS for both IPv4 and IPv6 then IPv4 and IPv6 topologies MUST match exactly. Cannot run IS-IS IPv6 on some interfaces, IS-IS IPv4 on others.
- IS-IS performs check for IPv4 address presence and matching subnet

```
Router(config-router)# no adjacency-check
```

- Adjacencies are formed only between matching setups eg. IS-IS IPv6 only router and IPv4 only router will not form adjacency.
  - For existing IPv4 only topology migrating to IPv6

```
Router(config-router)# address-family ipv6
Router(config-router-af)# no adjacency-check
```

# Summary

- Integrated IS-IS routing for IP uses CLNS and therefore requires CLNS addresses, that is, NET addresses.
- Integrated IS-IS requires planning the addresses, enabling the router, defining the router NET, and enabling the appropriate interfaces.
- IS-IS can be optimized by adjusting adjacency levels and changing the default metric cost.
- IS-IS summarization can be configured with the **`summary-address`** command.

# Additional Links

- http://www.cisco.com/en/US/customer/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800c5bc1.html
- http://www.certificationzone.com/cisco/newsletter/SL/IE-ISIS-WP2-F02_BIF.html
- http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcprt2/1cfisis.htm

# The Show Commands

- `show ip protocols`
- `show ip isis interfaces`
- `show ip route`

Slides adapted by Vladimír Veselý and Marcel Marek partially from official Cisco course materials but the most of the credit goes to CCIE#23527 Ing. Peter Palúch, Ph.D.

Last update: 2016-03-29