



*TOMORROW
starts here.*

MPLS
Workshop



ISIS Fundamentals and Troubleshooting

BRKRST-2327

MPLS
Workshop

ISIS Fundamentals and Troubleshooting

Agenda

- Overview
- Hierarchical Areas
- Addressing
- MTU and Hello Padding
- Attach-bit and Route Leaking
- LSP Flooding and Convergence
- SPF and Network Stability
- Route Redistribution
- Narrow and Wide Metrics



Reference slide



What Is IS-IS?

- Intermediate **S**ystem **to** Intermediate **S**ystem
- An “IS” is ISO terminology for a router
- IS-IS was originally designed as a dynamic routing protocol for ISO CLNP, defined in the ISO 10589 standard
- Also called as “Integrated IS-IS” or “Dual IS-IS”
- Encodes the packet(s) in TLV (Type, Length, Value) format



IS-IS for IP Routing

- Easily adapted to carry IP prefix information, as specified in RFC1195
- Flexible protocol in terms of tuning and easily extensible with to new features with TLVs
 - IS-IS extensions for MPLS -TE, IPv6
- It runs directly over Layer 2
- Proven to be a very stable and scalable, with very fast convergence

Encapsulation of IS-IS

- IS-IS is **not** encapsulated in IP!
- Encapsulated directly in the data link layer
- Protocol family is OSI: 0xFEFE is used in layer 2 headers to identify it

IS-IS:

Data-link header
(OSI family 0xFEFE)

IS-IS fixed header
(first byte is 0x83)

IS-IS TLVs



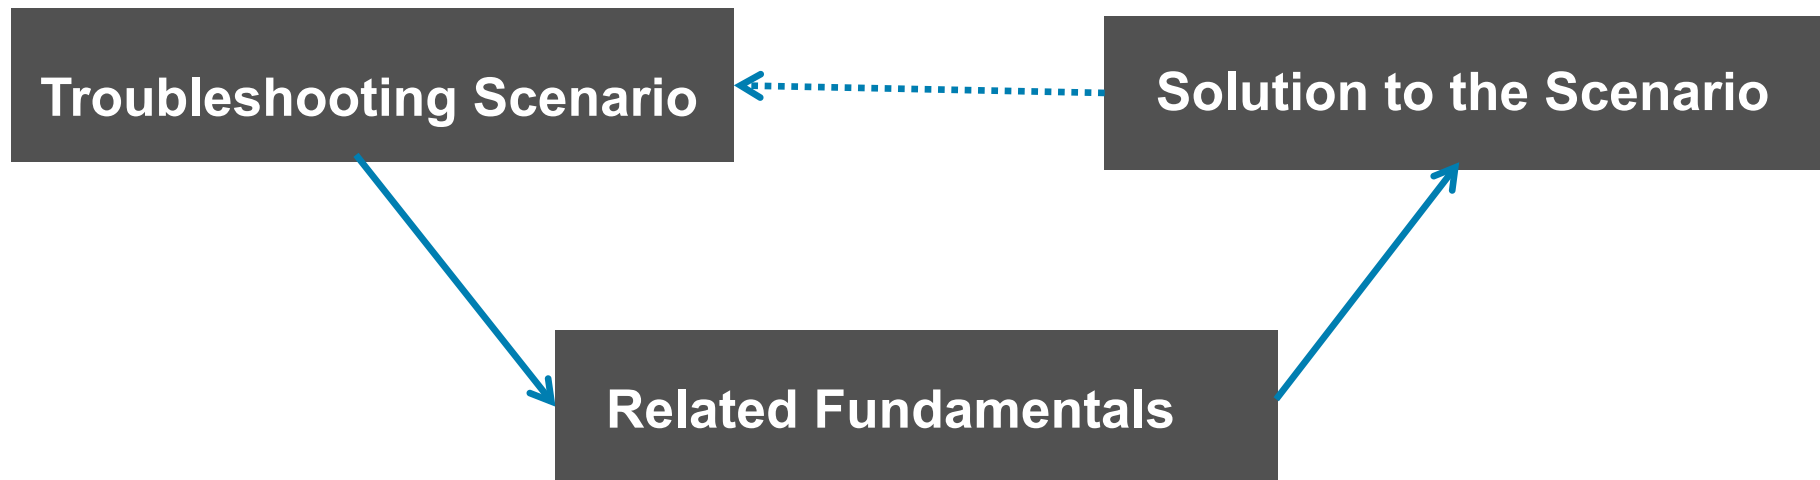
ISIS Fundamentals and Troubleshooting

Agenda

- Overview
- IS-IS Hierarchical Areas
- Addressing
- MTU and Hello Padding
- Attach-bit and Route Leaking
- LSP Flooding and Convergence
- SPF and Network Stability
- Route Redistribution
- Narrow and Wide Metrics


ISIS Fundamentals and Troubleshooting

Flow of the presentation



Scenario 1 :

R2 and R4 Peering showing as ES-IS instead of IS-IS?



R2#show clns neighbors

| System Id | Interface | SNPA | State | Holdtime | Type |
|-----------|-----------|----------------|-------|----------|------|
| Protocol | | | | | |
| R1 | Et0/0 | 00d0.58eb.ff01 | Up | 26 | L1 |
| R4 | Se3/0 | *HDLC* | Up | 280 | IS |

ES-IS
MPLS
Workshop



Hierarchy Levels

- IS-IS presently has a two-layer hierarchy
 - The backbone (level 2)
 - The areas (level 1)
- An IS (router) can be either:
 - Level 1 router (intra-area routing)
 - Level 2 router (inter-area routing)
 - Level 1–2 router (intra and inter-area routing)



Level 1 Routers

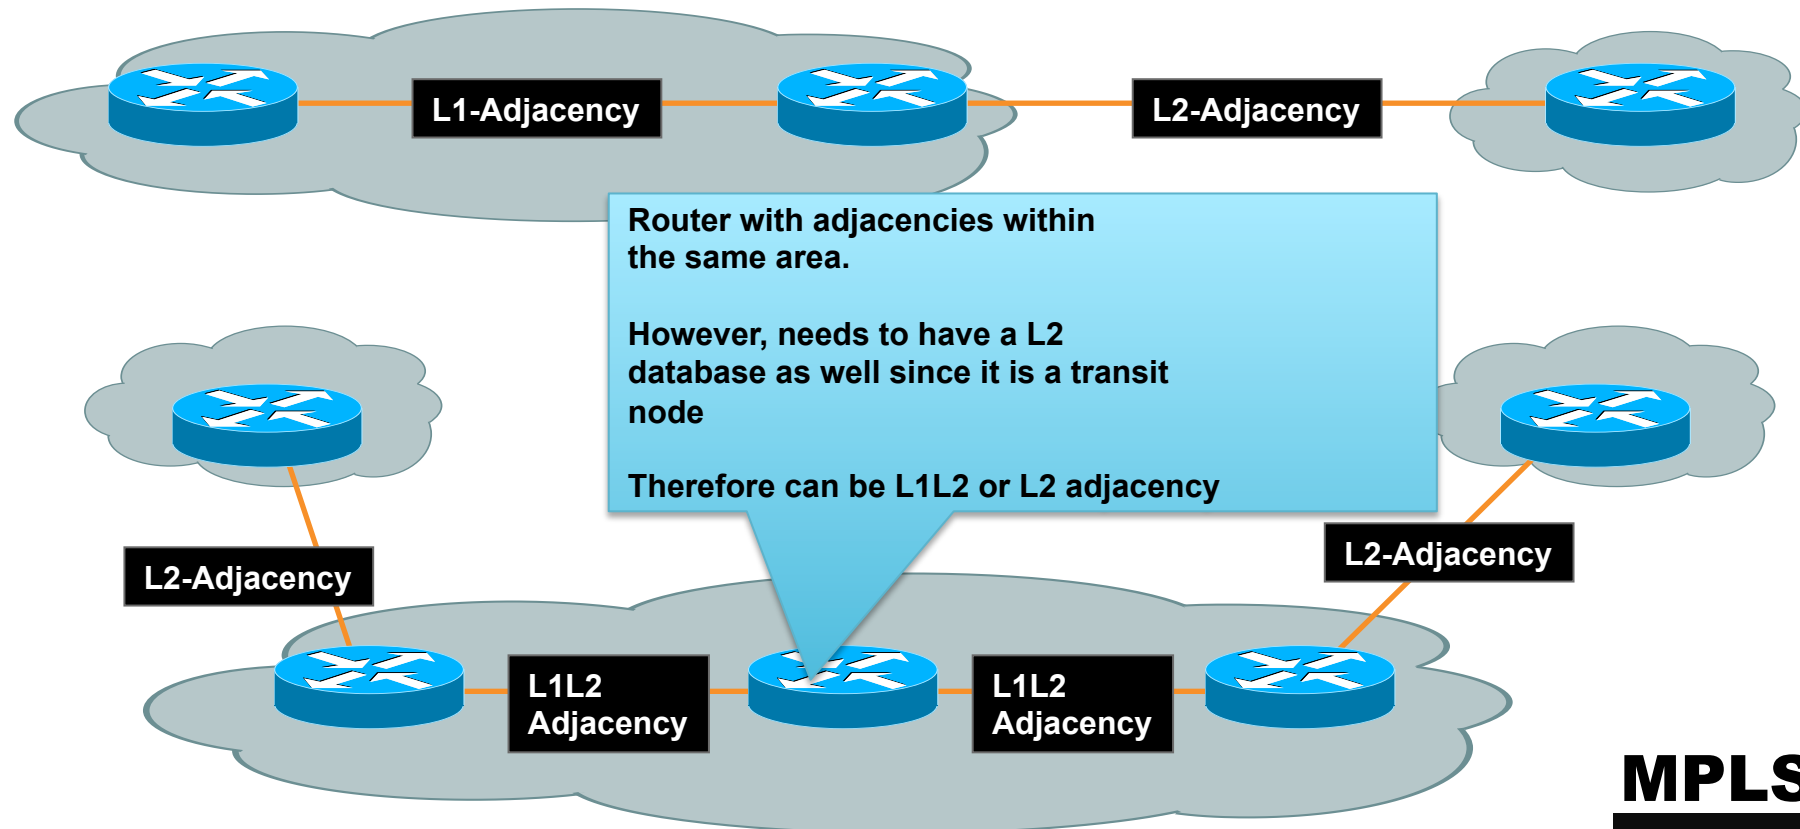
- Neighbors only in the **same** Level 1 area
- Level 1 LSDB only carries intra-area information
- Level 1 only routers look at the attached-bit to find the closest Level 1–2 router
- Level 1 only routers install a default route to the closest Level 1–2 router in the area



Level 2 Routers

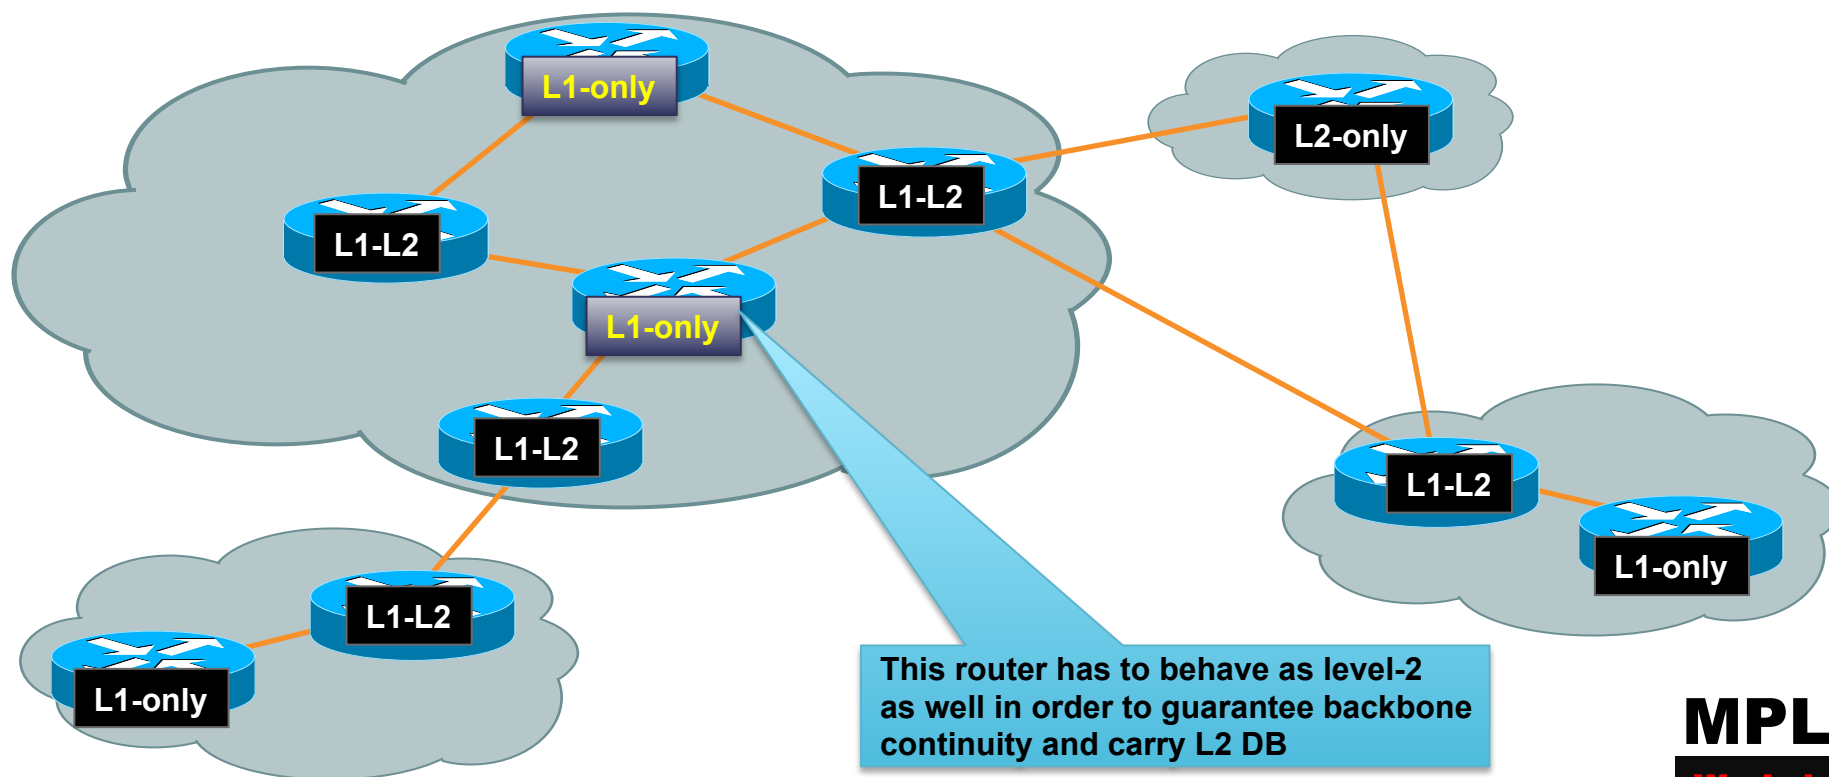
- May have neighbors in other areas
- L2 has information about L2 topology
- L2 has information about which L1 areas are reachable and how to reach them via the L2 topology
- L2 routers often may also perform L1 routing
 - called **L1L2** routers

Adjacency Levels

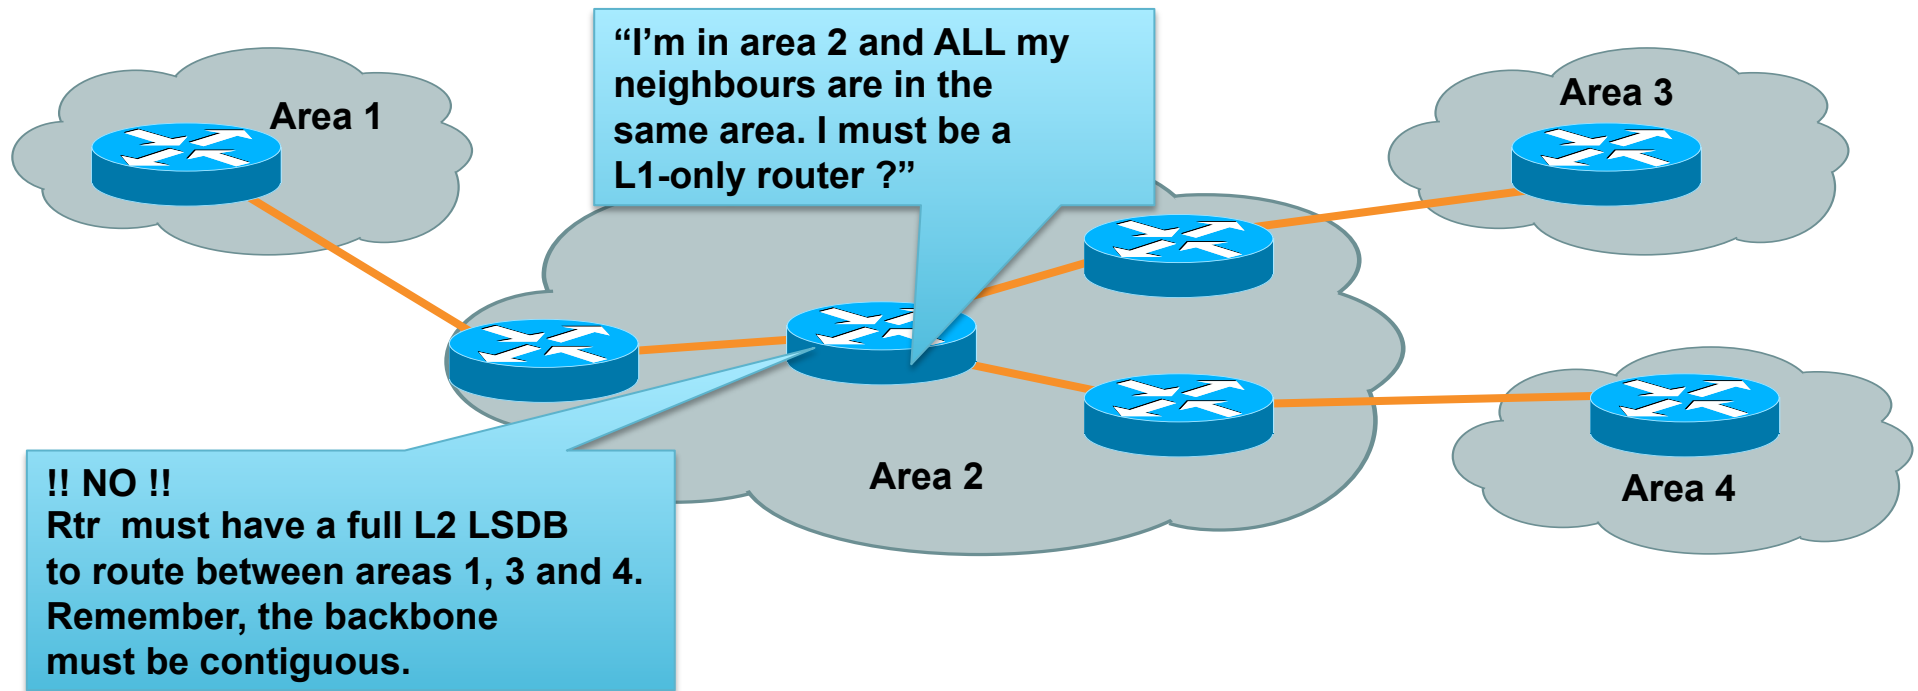


Level 1, Level 2 and Level 1–2 Routers

Backbone **Must Be** L2 Contiguous




Can an IS Determine Its Level?




- IS-IS router cannot determine if they need to be L1 or L2; therefore, by default all Cisco routers will behave as **L1L2**

Areas and Backbone Routers - Example




How IS-IS Area is different from OSPF ?

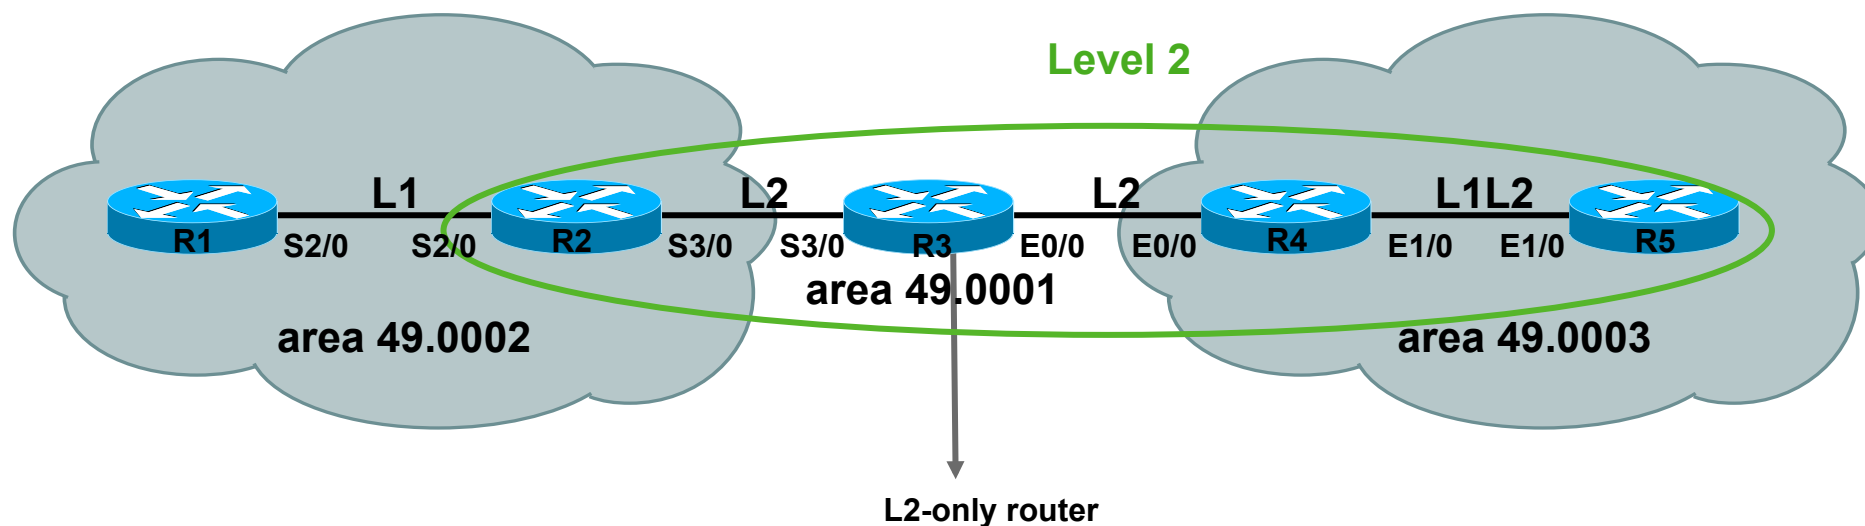


How IS-IS Area is different from OSPF ?

- OSPF allows dis-contiguous backbone, a virtual link can be used to bridge them
- There is no Virtual-link in IS-IS, L2 areas / Backbone must be contiguous




Basic IS-IS Configuration for IP Network




Area 49.0001 does not need to have level 1 enabled

Basic IS-IS Configuration for IP Network – R1




```
interface Loopback0
 ip address 10.1.100.1 255.255.255.255
!
interface Serial2/0
 ip address 10.1.5.1 255.255.255.0
 ip router isis
!
router isis
 net 49.0002.0000.0000.0001.00
 is-type level-1
 log-adjacency-changes all
 passive-interface Loopback0
```


Basic IS-IS Configuration for IP Network – R2




```
interface Serial2/0
 ip address 10.1.5.2 255.255.255.0
 ip router isis
 isis circuit-type level-1
!
interface Serial3/0
 ip address 10.1.1.2 255.255.255.0
 ip router isis
 isis circuit-type level-2-only
!
router isis
 net 49.0002.0000.0000.0002.00
 log-adjacency-changes all
```

Basic IS-IS Configuration for IP Network – R3




```
interface Ethernet0/0
 ip address 10.1.2.3 255.255.255.0
 ip router isis
!
interface Serial3/0
 ip address 10.1.1.3 255.255.255.0
 ip router isis
!
router isis
 net 49.0001.0000.0000.0003.00
 is-type level-2-only
 log-adjacency-changes all
```

Basic IS-IS Configuration for IP Network – R3



```
interface Ethernet0/0
 ip address 10.1.2.4 255.255.255.0
 ip router isis
 isis circuit-type level-2-only
!
interface Ethernet1/0
 ip address 10.1.4.4 255.255.255.0
 ip router isis
 isis network point-to-point
!
router isis
 net 49.0003.0000.0000.0004.00
 log-adjacency-changes all
 passive-interface Loopback0
```

Basic IS-IS Configuration for IP Network – R5




```
interface Loopback0
 ip address 10.1.100.5 255.255.255.255
!
interface Ethernet1/0
 ip address 10.1.4.5 255.255.255.0
 ip router isis
 isis network point-to-point
!
router isis
 net 49.0003.0000.0000.0005.00
 log-adjacency-changes all
 passive-interface Loopback0
```

network type point-to-point

Basic IS-IS Configuration for IP Network

'show clns'



Check the interfaces enabled and mode with 'show clns' command

R3#show clns

Global CLNS Information:

2 Interfaces Enabled for CLNS

NET: 49.0001.0000.0000.0003.00

Configuration Timer: 60, Default Holding Timer: 300, Packet Lifetime 64

ERPDU's requested on locally generated packets


Running IS-IS in IP-only mode (CLNS forwarding not allowed)

"clns routing" is not configured

MPLS
Workshop

Basic IS-IS Configuration for IP Network

'show clns protocol'



```
R3#show clns protocol
```

```
IS-IS Router: <Null Tag>
```

```
System Id: 0000.0000.0003.00 IS-Type: level-2
```

```
Routing for area address(es):
```

```
49.0001
```

```
Interfaces supported by IS-IS:
```

```
Ethernet0/0- IP
```


```
Serial3/0 - IP → no "OSI"
```

```
...
```

```
Generate narrow metrics: level-1-2
```

```
Accept narrow metrics: level-1-2
```


Solution to Scenario 1 – Session showing ES-IS



```
R2#show clns neighbors
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|--------------|
| R1 | Et0/0 | 00d0.58eb.ff01 | Up | 26 | L1 | IS-IS |
| R4 | Se0/0 | *HDLC* | Up | 280 | IS | ES-IS |

MPLS
Workshop




ISIS Fundamentals and Troubleshooting

Agenda

- Overview
- IS-IS Hierarchical Areas
- Addressing
- MTU and Hello Padding
- Attach-bit and Route Leaking
- LSP Flooding and Convergence
- SPF and Network Stability
- Route Redistribution
- Narrow and Wide Metrics

Scenario 2 : R1 sees other routers only as L2 peers



R1#show clns neighbors

| System Id | SNPA | Interface | State | Holdtime | Type | Protocol |
|-----------|----------------|-----------|-------|----------|------|----------|
| R2 | 0000.0c76.f098 | Et0 | Up | 27 | L2 | IS-IS |
| R3 | 0000.0c76.f096 | Et0 | Up | 26 | L2 | IS-IS |




Integrated IS-IS: Addressing

- The address at which the network service is accessible is known as the **NSAP**:
 - **Network Service Access Point**
- One NSAP per router, not per interface
- An NSAP can be a total of 20 bytes long

Basic NSAP Format

- An NSAP mainly consists of three parts:



- Total length is between eight and 20 bytes
- Example:

| | | |
|--------|-----------------|----|
| 49.01. | 0000.0000.0007. | 00 |
|--------|-----------------|----|

Area System ID N-SEL (always zero for a Router)



NETs versus NSAPs

- **NET**: Network Entity Title
 - Is the address of the network entity itself
- A **NET** is an NSAP where N-selector is 0 (common practice)
 - A **NET** implies the routing layer of the IS itself (no transport layer)
- ISs (routers) do not have any transport layer (so, always selector=0). N-Sel is a Non-zero value only for Pseudonodes
- Multiple **NETs** are like secondary IP addresses
 - Proper use is to only use them when merging or splitting areas



NET Address Examples

Total length of NET can be between 8 and 20 bytes

- Example 1:

- 47.0001.aaaa.bbbb.cccc.00

Area = 47.0001

SysID = aaaa.bbbb.cccc

NSel = 00

- Example 2:

- 39.0f01.0002.0000.0c00.1111.00

Area = 39.0f01.0002

SysID = 0000.0c00.1111

NSel = 00

- Example 3:


- 49.0002.0000.0000.0007.00

Area = 49.0002

SysID = 0000.0000.0007

Nsel = 00

CLNS Addressing: Example



Creating Unique SystemIDs



- SystemID is 6 bytes
- Some methods to create unique SystemIDs:
 - Start numbering 1, 2, 3, 4 etc
 - Use MAC addresses
 - Convert a loopback IP address
 - **192.31.231.16 -> 192.031.231.016 -> 1920.3123.1016 -> systemID**

One Way of Defining NSAPS

- Take the loopback IP address of the router and make it SystemID
- Define the area
 - Process:

Take the loopback IP address

**Fill up with zero's to reach three digits
between dots**

**Move dots to have three groups of 4
digits to form address**

Merge area and address

10.1.3.120

010.001.003.120

0100.0100.3120

49.0123.0100.0100.3120

MPLS

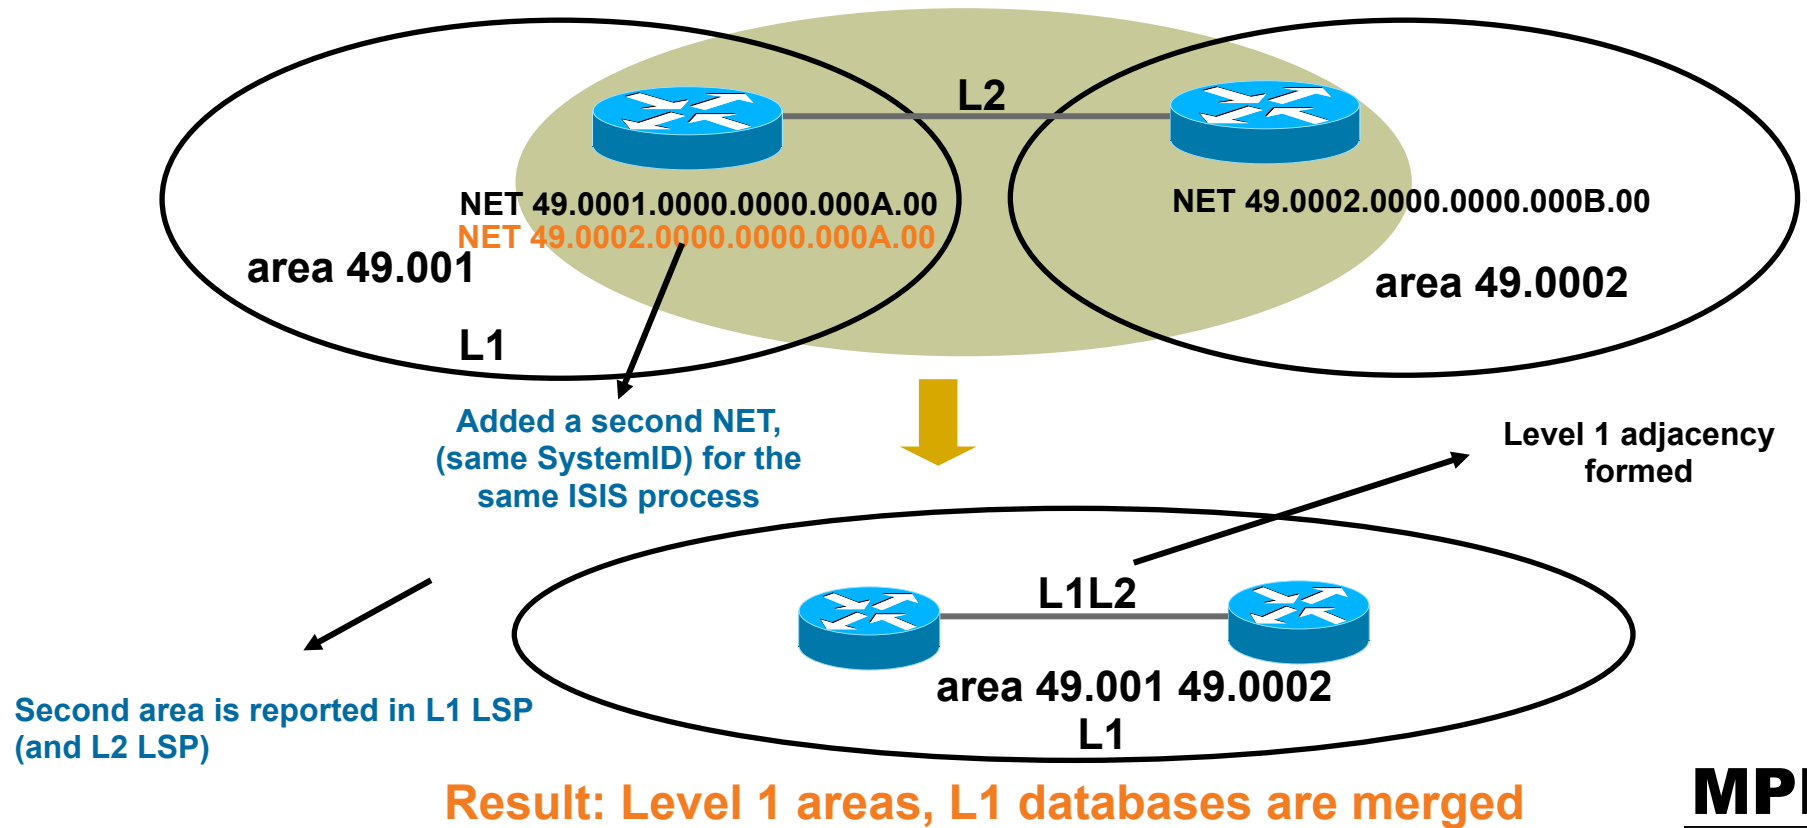
Workshop

Multiple NETs




- Configuring multiple NETs on one ISIS router leads to merging those areas
 - The result is one level-1 area
 - All NETs must have the same System ID
 - Otherwise each router would originate multiple LSPs
- Up to three NETs can be configured in IOS
- Reasons:
 - Multiple NETs can be used to merge areas in the transition period
 - Multiple NETs can be used to split areas in the transition period
 - Can be used to change the area address
 - More than one addressing authority for one area
- Two L1 ISIS routers become adjacent if they share at least one area in the different NETs

Multiple NETs: Example



Duplicate System id



Duplicate System ID Errors in logs :

```
R1#show logging | i Duplicate
```


```
Apr 9 16:41:20: %CLNS-3-BADPACKET: ISIS: LAN L1 hello, Duplicate system ID det)
Apr 9 16:42:22: %CLNS-3-BADPACKET: ISIS: LAN L1 hello, Duplicate system ID det)
Apr 9 16:43:21: %CLNS-3-BADPACKET: ISIS: LAN L1 hello, Duplicate system ID det)
```

Duplicate System ID in debugs :

```
R1#debug isis adj-packet
```


```
Apr 9 16:41:53: ISIS-Adj: Sending L1 IIH on Ethernet0/0, length 1497
Apr 9 16:41:55: ISIS-Adj: Rec L1 IIH from 00d0.58eb.ff01 Ethernet0/0)
Apr 9 16:41:55: ISIS-Adj: Duplicate system id
```

IS-IS SystemID and OSPF RouterID



- One systemID / RouterID per router
- Globally unique across areas

Solution to Scenario 2 : R1 sees other routers as 'L2 only' !



R1#show clns neighbors

| System Id | SNPA | Interface | State | Holdtime | Type | Protocol |
|-----------|----------------|-----------|-------|----------|------|----------|
| R2 | 0000.0c76.f098 | Et0 | Up | 27 | L2 | IS-IS |
| R3 | 0000.0c76.f096 | Et0 | Up | 26 | L2 | IS-IS |

MPLS
Workshop




ISIS Fundamentals and Troubleshooting

Agenda

- Overview
- IS-IS Hierarchical Areas
- Addressing
- MTU and Hello Padding
- Attach-bit and Route Leaking
- LSP Flooding and Convergence
- SPF and Network Stability
- Route Redistribution
- Narrow and Wide Metrics

Scenario 3 : One way IS-IS, other way ES-IS



```
R2#show clns neighbors
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|--------|-------|----------|------|--------------|
| R4 | Se0/0 | *HDLC* | Up | 280 | IS | ES-IS |

```
R4#show clns neighbors
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|--------|-------|----------|------|--------------|
| R2 | Se0/0 | *HDLC* | Init | 27 | L2 | IS-IS |

MPLS
Workshop



Hello Padding

- IS-IS by default pads the Hellos to the full interface MTU size to detect MTU mismatches
- Useful to detect misconfigurations or underlying layer problems
 - Example: EoMPLS scenarios: a link over AToM might have MTU of 4k on the edges, but a lower MTU in the MPLS core
- If the operator is sure of the MTU on the link, the padding of the Hellos can be turned off
 - Avoid using bandwidth unnecessary
 - Reduced Buffer Usage
 - Reduced processing overhead when using authentication



Hello Padding

- Turn on/off the Hello-Padding either per interface level or globally
- Two ways of disabling :

Under router isis CLI

[no] hello padding [multi-point|point-to-point]

Under interface CLI

[no] isis hello padding

- Even if padding is disabled, at the beginning, the router still sends a few hellos at full MTU. 'always' option which is hidden can be used to prevent it.

Hello Padding – Check MTU with ‘show clns interface’

```
R1#show clns interface Ethernet 0/0
```

```
Ethernet0/0 is up, line protocol is up
```

```
Checksums enabled, MTU 1497, Encapsulation SAP
```

```
ERPDUs enabled, min. interval 10 msec.
```

```
Routing Protocol: IS-IS
```

```
Circuit Type: level-1-2
```

```
...
```

```
Number of active level-1 adjacencies: 1, if state UP
```

```
Level-2 Metric: 10, Priority: 64, Circuit ID: R1.01
```

```
Level-2 IPv6 Metric: 10
```

```
Number of active level-2 adjacencies: 1, if state UP
```


```
Next IS-IS Hello in 1 seconds
```

```
if state UP
```

CLNP, ES-IS, IS-IS use SAP encapsulation -> MTU = 1497

DSAP: 1 byte
SSAP: 1 byte
Control field: 1 byte

Solution to Scenario 3 : One way IS-IS, other way ES-IS



R2#show clns neighbors

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|--------|-------|----------|------|--------------|
| R4 | Se0/0 | *HDLC* | Up | 280 | IS | ES-IS |

R4#show clns neighbors

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|--------|-------------|----------|------|--------------|
| R2 | Se0/0 | *HDLC* | Init | 27 | L2 | IS-IS |

MPLS
Workshop



Reason for one way ES-IS

- R4 receives and processes the smaller 1499-byte hellos from R2 and puts the IS-IS adjacency in 'init' state, hoping to complete the three-way handshake to establish full IS-IS adjacency.

R4#debug isis adj-packet

```
Apr 9 20:44:16: ISIS-Adj: Sending serial IIH on Serial0/0, length 1999
Apr 9 20:44:21: ISIS-Adj: Rec serial IIH from *HDLC* (Serial0/0)
Apr 9 20:44:21: ISIS-Adj: rcvd state DOWN, old state UP, new state INIT
( Moved to INIT after peer's hello received )
Apr 9 20:44:21: ISIS-Adj: Action = GOING DOWN
```



Reason for one way ES-IS

- After three hellos are ignored by R2's Interface due to higher size, the hello hold time expires on R2, the adjacency is dropped, and an adjacency change event is logged. ES-IS does not pad hellos !

R2#debug isis adj-packet

```
Apr 9 20:43:56: ISIS-Adj: Sending serial IIH on Serial0/0, length 1499
Apr 9 20:44:05: ISIS-Adj: Sending serial IIH on Serial0/0, length 1499
Apr 9 20:44:13: ISIS-Adj: Sending serial IIH on Serial0/0, length 1499
Apr 9 20:44:22: ISIS-Adj: Sending serial IIH on Serial0/0, length 1499
( no hellos received as the peer is sending 2000 byte )
Apr 9 20:44:29: %CLNS-5-ADJCHANGE: ISIS: Adjacency to RT2 (Serial0/0)
Down
( Tearing down ISIS peering )
```




ISIS Fundamentals and Troubleshooting

Agenda

- Overview
- IS-IS Hierarchical Areas
- Addressing
- MTU and Hello Padding
- Attach-bit and Route Leaking
- LSP Flooding and Convergence
- SPF and Network Stability
- Route Redistribution
- Narrow and Wide Metrics

Scenario 4 :

Rtr A in Area 1 taking the Longer path to Rtr B in Area 2





Level-1 Routing

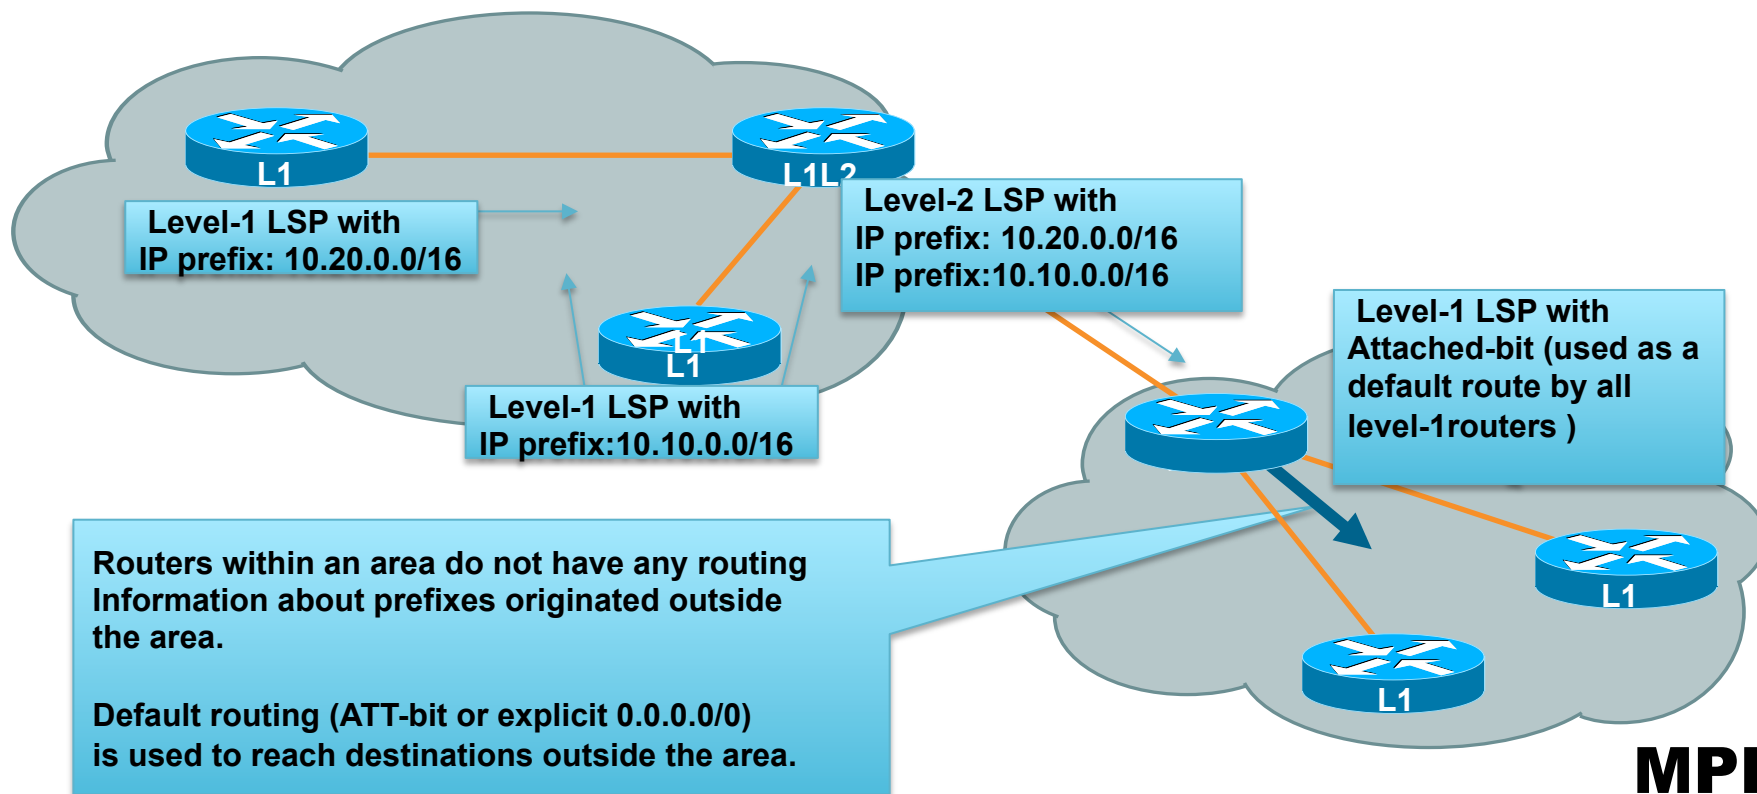
- L1-only routers know only topology of their own area
- Traffic to other areas is sent via the closest L2 IS
 - Can result in suboptimal routing
- L1L2 ISs set the “attached-bit” in their L1 LSP
- L1-only routers look at the attached-bit in L1 LSPs to find the closest L1L2 router
- L1-only routers install a default route to the closest L1L2 router in the area




L1 Advertised into L2

- All L1L2 routers advertise all the IP prefixes they learn via L1 into L2
- Only advertise routes you use
 - Inter-level routing goes via the RIB. In other words, If it is not in the routing table, it is not advertised from L1 to L2
- Summarization :
 - possible at L1->L2 or when redistributing
 - Internal ISIS routes are summarized only on L1-2 (from L1 to L2). External routes can be summarized even within the L1 area.
 - All level 1-2 routers in an area must summarize equivalently into the backbone to avoid traffic being sent to the same router (based on longest match)

IS-IS Routing Levels



The 'Attached' Bit



- L1L2 routers set the ATT bit in their L1 LSP
- L1 routers use ATT bit found in L1-LSDB as possible area exit point
 - ISIS for IP: level-1 router will install a 0.0.0.0/0 route toward the L1L2 with ATT-bit set
- Shortest metric to the L1L2 who sets the ATT bit wins



Route Leaking (L2 to L1)

- New ISIS feature/capability described in draft-ietf-isis-domain-wide
- Allows L1L2 routers to insert in their L1 LSP IP prefixes learned from L2 database if also present in the routing table
- ISIS areas are not stubby anymore



Route Leaking - Solution for Several Issues

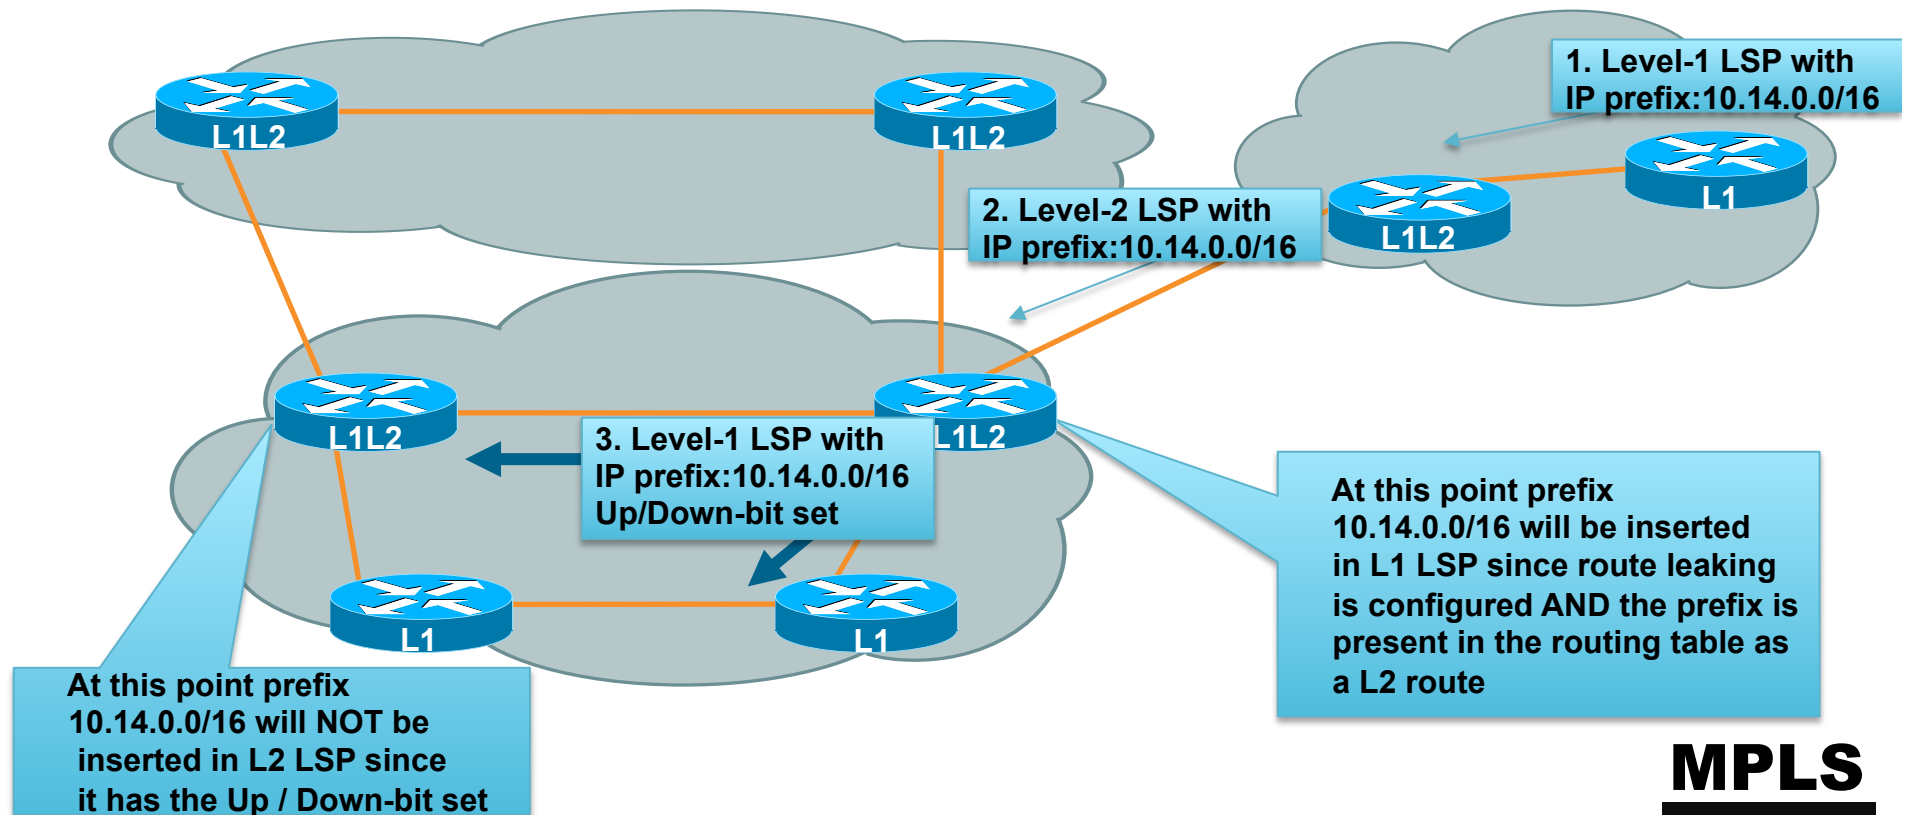
- Optimal inter-area routing
- BGP shortest path to AS exit point
- MPLS-VPN (PEs loopback reachability)



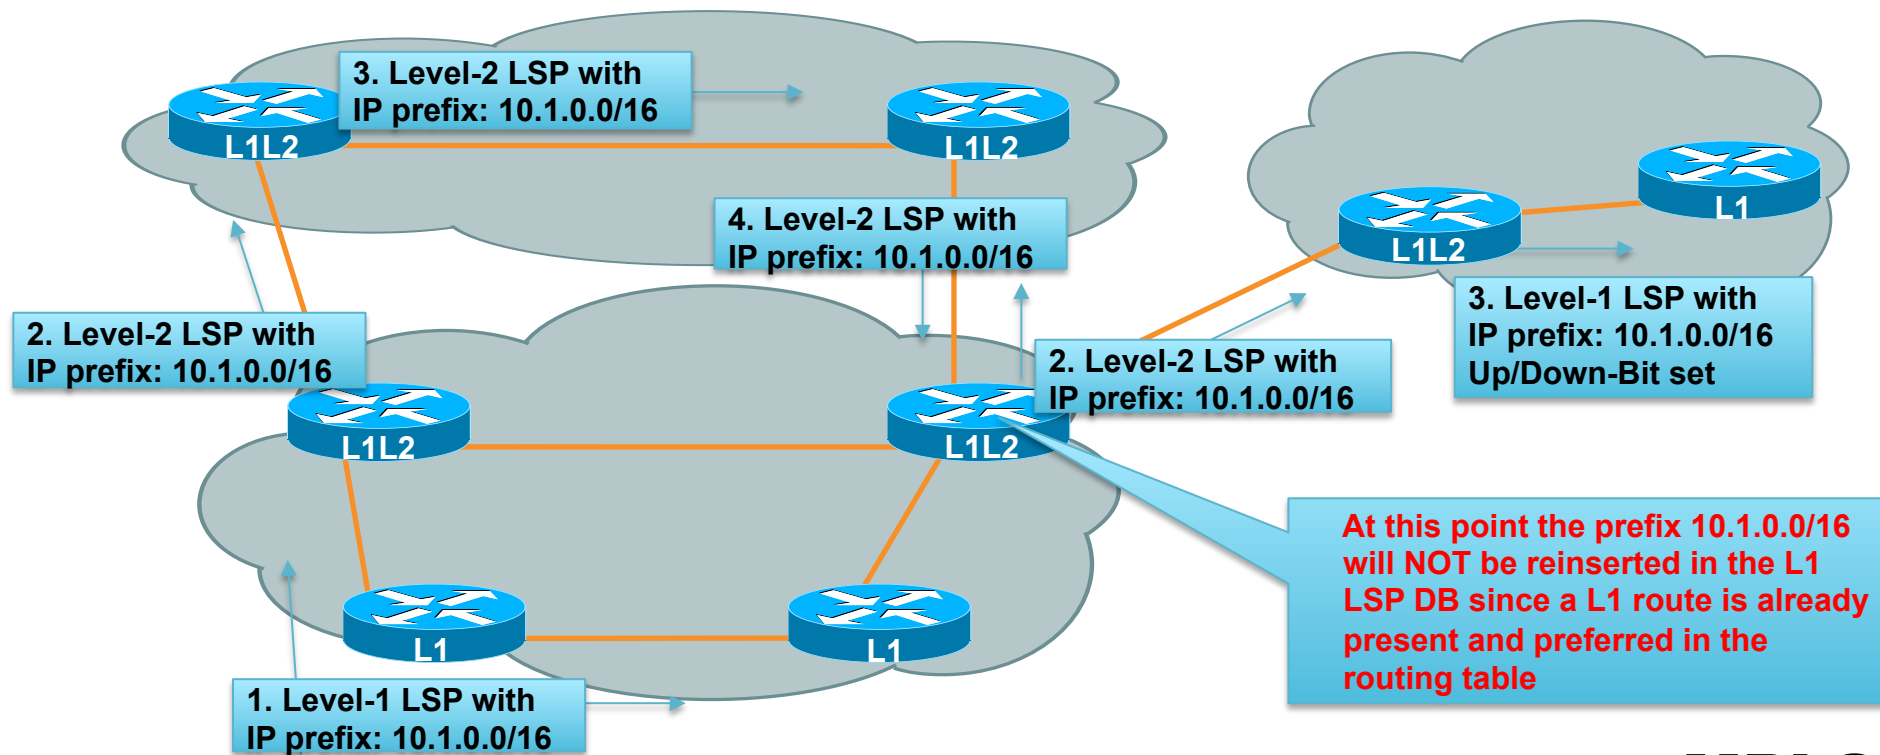
Route Leaking – Loop avoidance

- When leaking routes from L2 backbone into L1 areas a loop protection mechanism needs to be used in order to prevent leaked routes to be re-injected into the backbone
- UP/Down bit
 - Extended IP Reachability TLV (135) contains Up/Down bit
- UP/Down bit is set each time a prefix is leaked into a lower level
- Prefixes with Up/Down bit set are NEVER propagated to a upper level

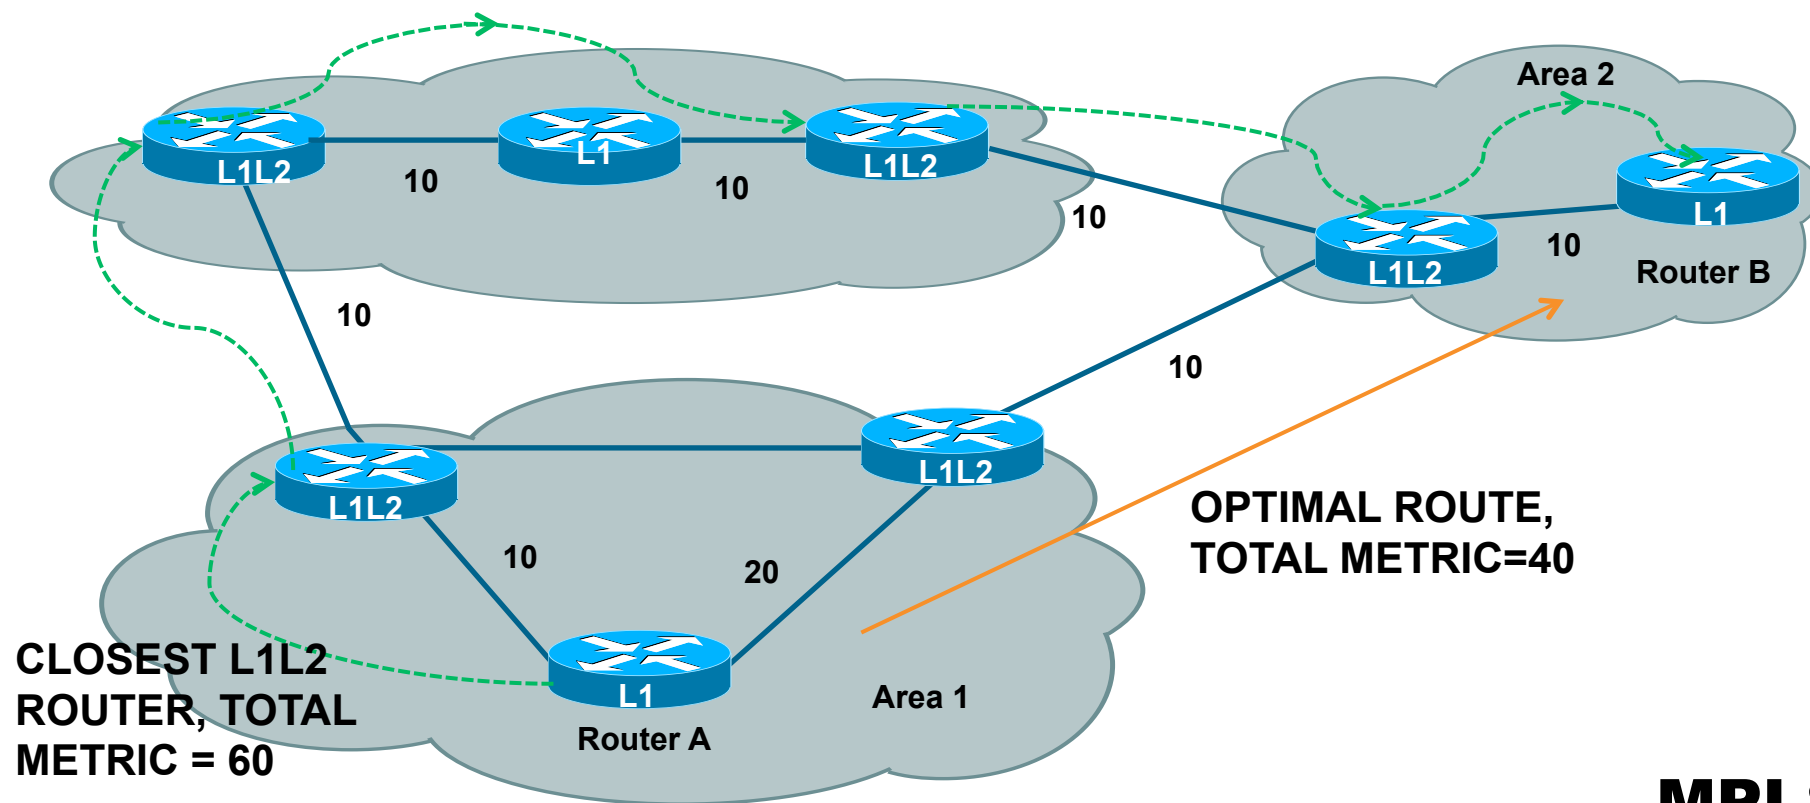
Route Leaking and Up/Down-bit



Route Leaking – Reinsertion prevented by default



Solution to Scenario 4 : Why Longer path ?



Configuring Route Leaking

On the L1-L2 Routers :

```
RtrA#(conf)router isis
RtrA#(conf-router)redistribute isis ip level-2 into level-1
distribute-list <100-199>
```

A distribute-list could be used if we want to leak only a set of routes


On L1 router, the routes leaked by L1-L2 Router are seen as Inter-Area:

RtrA# show ip route

i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area * - candidate default,

```
i ia 1.0.0.0/8 [115/30] via 55.55.55.1, Serial1/0
i ia 2.0.0.0/8 [115/30] via 55.55.55.1, Serial1/0
```

Scenario 5 : L1 to L2 route not leaked



```
R3#sh ip route 88.88.88.88
% Network not in table
```

```
R3#sh ip route 99.99.99.99
Known via "isis", distance 115, metric 30, type level-2
* 172.16.2.2, from 172.16.2.2, 00:10:07 ago, via Ethernet0/0
```




Inter-level routing goes via the RIB

- Though L1 to L2 leaking is done by default, routes need to be in the RIB as ISIS route for the leaking to succeed on the L1 / L2 Router
- L1 / L2 router advertises only the routes it uses
- The same holds true for L2 to L1 as well, though leaking needs to be manually configured in this case

Solution to Scenario 5 :

Why 88.88.88.88 is missing on R3 ?

Not present in R2's RIB as ISIS, so
not leaked from L1 to L2 on R2



```
R2#show ip route 88.88.88.88
```

```
Routing entry for 3.3.3.3/32
```

```
Known via "ospf 1" , distance 110, metric 2, type intra area
```



ISIS Fundamentals and Troubleshooting

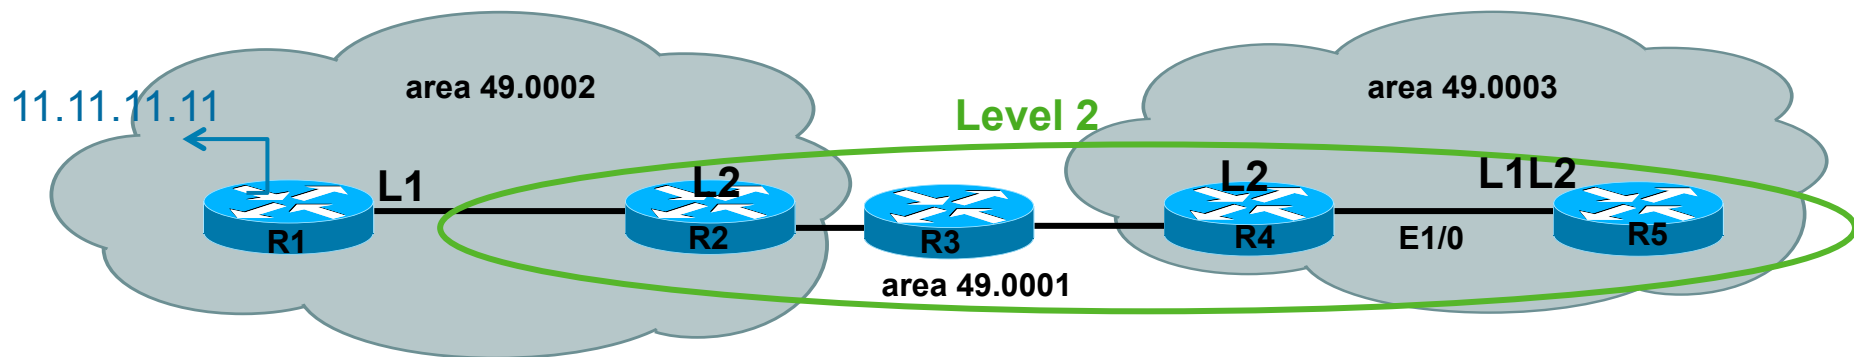
Agenda

- Overview
- IS-IS Hierarchical Areas
- Addressing
- MTU and Hello Padding
- Attach-bit and Route Leaking
- LSP Flooding and Convergence
- SPF and Network Stability
- Route Redistribution
- Narrow and Wide Metrics

Scenario 6 :

R5 does not have routes to Networks behind R1

LSP Info not reaching



```
R5#sh ip route 11.11.11.11
% Network not in table
```

MPLS
Workshop



IS-IS Packet Types

- Hello Packets (IIH)
 - Used to form adjacencies
- Link State Packets (LSP)
 - Describes the state of each router
- Sequence Number Packets (SNP)
 - PSNPs: Used for acknowledgements
 - CSNPs: Used to describe the LSPDB
- In ISO terminology, packets are referred to as Protocol Data Units (PDUs)

IS-IS Hello Packets



- Also called IIHs
- Used for maintaining adjacencies
- Different on point-to-point links and LANs
- By default, IIHs are padded to full MTU size
 - To maintain link integrity
 - Padding can be removed by configuration



Link State Packets

- Also known as Link State PDUs
- Contains all information about one router
- One (set of) LSP(s) per router
- One (set of) LSP(s) per LAN network



Sequence Number Packets

- Used when flooding the LSDB. Also known as Sequence Number PDUs.
- Two Types of SNPs
 1. Partial Sequence Number Packets (PSNPs)
 2. Complete Sequence Number Packets (CSNPs)
- PSNPs are like ACKs on point-to-point links
- CSNPs are used for LSDB sync over LANs
- CSNPs are also used to sync LSDB over new point-to-point (p2p) adjacent

MPLS

Workshop




What is a Pseudonode ?


- For SPF, the whole network must look like a collection of nodes and p2p links, however Multi-access networks are different.
- Create a virtual node, or **Pseudonode**
- Not a real router, but extra LSP in IS-IS
- Allows for smaller, more stable LSPs
- Similar to 'Network LSA' created by OSPF DR

Pseudonode on a LAN


Physical View



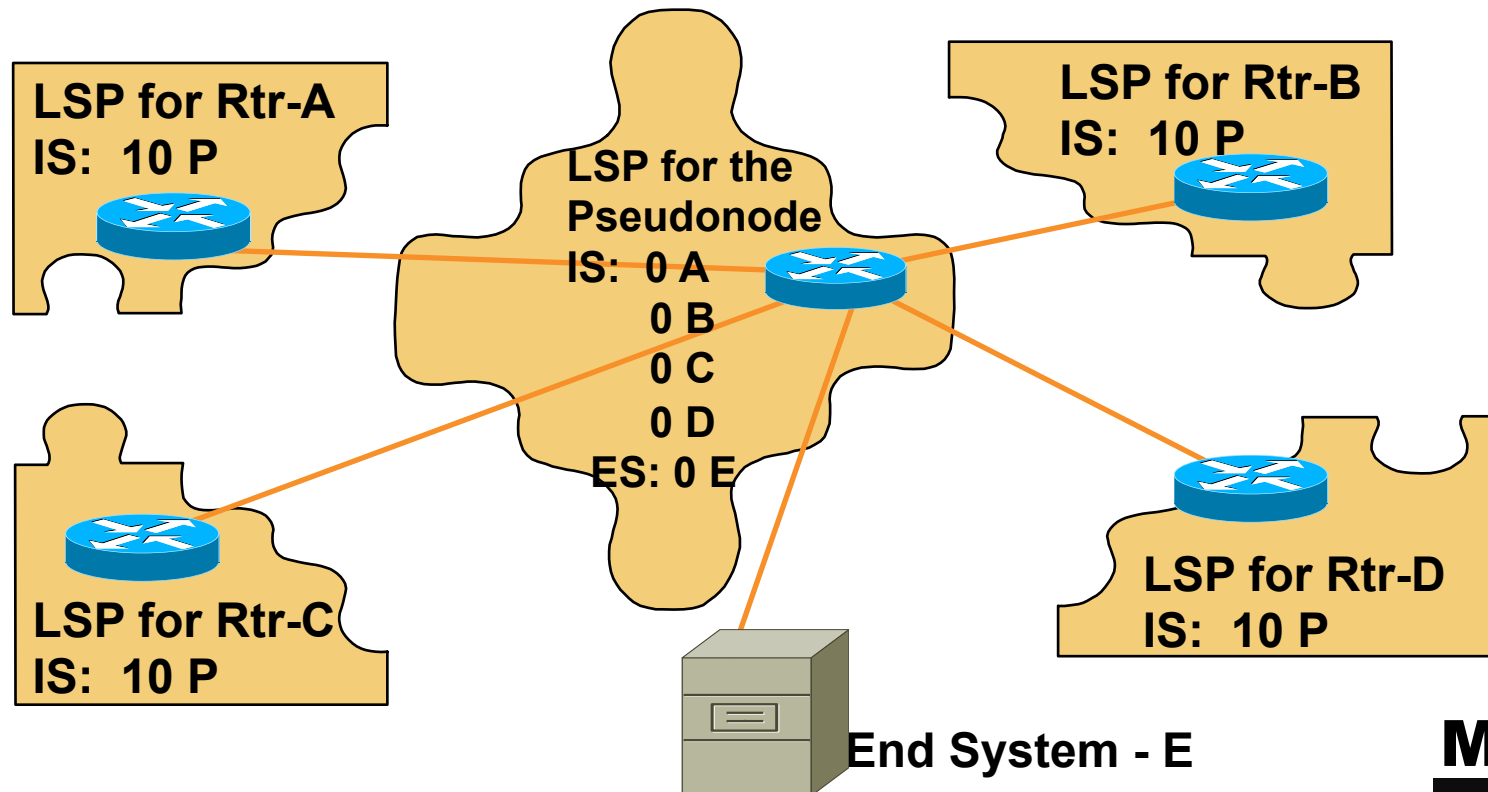
Logical View



LSDB without Pseudonode



LSDB with Pseudonode





Who Creates the Pseudonode?

- Created by Designated Router (DIS)
- The DIS reports all LAN neighbors in the pseudonode LSP
- All LAN routers report connectivity to the pseudonode in their LSPs



ISIS DIS

- On broadcast multi-access networks, a single router is elected as the DIS
- In a segment, one of the routers elects itself the DIS, based on interface priority (the default is 64)
 - SNPA is compared if the interface priority is the same for all routers
 - SNPA is the mac address in LAN and DLCI in case of frame relay
 - System-id acts as a tie-breaker if SNPA is the same as well (in DLCI scenario)



ISIS DIS versus OSPF DR

- Unlike OSPF, there is no Backup DIS elected in ISIS
- Unlike OSPF, the DIS election is **preemptive**
 - If a new router boots on the LAN with a higher interface priority, the new router becomes the DIS
 - It purges the old pseudonode LSP and floods a new set of LSPs
 - one unstable router can potentially make the network unstable with flooding

What Triggers a New LSP?



When Something Changes...

- Adjacency came up or went down
- Interface up/down with IP prefix
- Redistributed IP routes change
- Inter-area IP routes change
- An interface is assigned a new metric



Basic Flooding Rules

When Receiving a LSP, Compare with LSP in LSDB :

- verify that checksum is correct
- If LSP received is newer, install it in the LSDB and flood to **all other neighbors** and then check if you need to run SPF
- If older, send newer LSP from our LSDB
- If we have same LSP, only send an 'ack'



Types of Flooding

- Flooding on p2p links with positive acks
 - Each LSP is acknowledged with a PSNP
- Flooding on LANs with negative acks
 - DIS multicasts a full list of LSP descriptions in a CSNP packet
 - Re-transmission requests are done via PSNP
- General background flooding

Flooding on a LAN



- LAN flooding is reliable due to the DIS
 - Creates and updates Pseudonode LSP
- DIS broadcasts CSNPs every 10 seconds
 - Negative acks
 - No backup DIS



Remaining Lifetime

- Used to age out old LSPs
- Periodic refresh needed to keep stable LSPs valid
- IS-IS counts down from 1200 seconds to 0
 - Time is configurable
- When lifetime expires, the LSP is purged from the network

Remaining Lifetime: Purging LSPs




- When Remaining Lifetime is zero, LSP is purged
- Detecting router removes LSP body and floods the header with RL=0
- All other routers then remove this LSP from their database
- On LANs, a new DIS purges pseudonode LSP

LSP Refresh



- Specifies the number of seconds a router will wait before refreshing its own LSP
- Only the originating router can re-create and re-flood its own LSP
- Can cause unnecessary overhead and limit scalability
- Default refresh is 15 minutes

Reading ISIS Database



R4#show isis database

* means "this router"

IS-IS **Level-1** Link State Database:

| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|----------|--------------|--------------|--------------|----------|
| R4.00-00 | * 0x0000007F | 0x27D1 | 762 | 1/0/0 |
| R5.00-00 | 0x00000076 | 0x42BC | 760 | 1/0/0 |


Attach bit is set on L1 LSPs from R4 and R5

IS-IS **Level-2** Link State Database:

| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|----------|--------------|--------------|--------------|----------|
| R2.00-00 | 0x0000007C | 0x548D | 814 | 0/0/0 |
| R3.00-00 | 0x00000081 | 0x28A1 | 820 | 0/0/0 |
| R4.00-00 | * 0x0000007A | 0xFF38 | 761 | 0/0/0 |
| R4.02-00 | * 0x00000065 | 0x7CF1 | 758 | 0/0/0 |
| R5.00-00 | 0x0000000A | 0x6CFB | 757 | 0/0/0 |

pseudonode LSP

Reading ISIS Database




R3#show isis database

IS-IS Level-2 Link State Database:

| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|----------|--------------|--------------|--------------|----------|
| R2.00-00 | 0x0000007D | 0x528E | 1022 | 0/0/0 |
| R3.00-00 | * 0x00000082 | 0x26A2 | 952 | 0/0/0 |
| R4.00-00 | 0x0000007B | 0xFD39 | 925 | 0/0/0 |
| R4.02-00 | 0x00000066 | 0x7AF2 | 975 | 0/0/0 |
| R5.00-00 | 0x0000000B | 0x6AFC | 938 | 0/0/0 |

(Please note R3 only has L2 LSP database)

Reading ISIS Database



R1#show isis database

IS-IS **Level-1** Link State Database:


| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|----------|--------------|--------------|--------------|---------------|
| R1.00-00 | * 0x0000006F | 0xD44D | 664 | 0/0/0 |
| R2.00-00 | 0x0000007D | 0xCB3B | 1154 | 1 /0/0 |

**Attach bit is set on L1
LSP from R2**

(Please note R1 only has L1 LSP database)

MPLS
Workshop

ISIS Database Detail



```
R4#show isis database R4.00-00 detail
```

```
IS-IS Level-1 LSP R4.00-00
```

| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|----------|--------------|--------------|--------------|----------|
| R4.00-00 | * 0x00000080 | 0x25D2 | 783 | 1/0/0 |

```
Area Address: 49.0003
```

```
NLPID: 0xCC
```

```
Hostname: R4
```


```
IP Address: 10.1.100.4
```

```
Metric: 10 IP 10.1.4.0 255.255.255.0
```

```
Metric: 0 IP 10.1.100.4 255.255.255.255
```

```
Metric: 10 IS R5.00 ( continued on next slide )
```


ISIS Database Detail - continued



IS-IS Level-2 LSP R4.00-00

| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|------------------------|-------------------------------|--------------|--------------|----------|
| R4.00-00 | * 0x0000007B | 0xFD39 | 666 | 0/0/0 |
| Area Address: 49.0003 | | | | |
| NLPID: 0xCC | | | | |
| Hostname: R4 | | | | |
| IP Address: 10.1.100.4 | | | | |
| Metric: 10 | IP 10.1.2.0 255.255.255.0 | | | |
| Metric: 10 | IS R4.02 | | | |
| Metric: 10 | IS R5.00 | | | |
| Metric: 10 | IP 10.1.4.0 255.255.255.0 | | | |
| Metric: 10 | IP 10.1.100.5 255.255.255.255 | | | |
| Metric: 0 | IP 10.1.100.4 255.255.255.255 | | | |

ISIS Pseudonode-LSP Example



R4#show isis database level-2

IS-IS Level-2 Link State Database:

| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|----------|-------------|--------------|--------------|----------|
| R2.00-00 | 0x00000003 | 0xDA80 | 954 | 0/0/0 |
| R4.01-00 | *0x00000001 | 0x4C87 | 954 | 0/0/0 |
| R5.00-00 | 0x00000003 | 0x0E61 | 956 | 0/0/0 |


R4#show isis database level-2 R4.01-00 detail

IS-IS Level-2 LSP R4.01-00

| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|-----------|-------------|--------------|--------------|----------|
| R4.01-00 | 0x00000001 | 0x4C87 | 914 | 0/0/0 |
| Metric: 0 | IS R4.00 | | | |
| Metric: 0 | IS R3.00 | | | |

IS neighbors of the DIS: R4 and R3 (L2!)

ISIS Fragmented LSP



```
R5#show isis database level-1
```


IS-IS Level-1 Link State Database:

| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|----------|--------------|--------------|--------------|----------|
| R4.00-00 | 0x00000019 | 0xF36B | 1005 | 1/0/0 |
| R5.00-00 | * 0x0000001F | 0xCF9F | 1100 | 1/0/0 |
| R5.00-01 | * 0x00000002 | 0x5BD7 | 1100 | 0/0/0 |

Two fragments of R5.00-00

Segment 0 always contains the critical data

ISIS Fragmented LSP – First Segment



```
R5#show isis database level-1 R5.00-00 detail
```

```
IS-IS Level-1 LSP R5.00-00
```

| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|------------|-------------|--------------|--------------|----------|
| R5.00-00 * | 0x0000001F | 0xCF9F | 976 | 1/0/0 |

```
Area Address: 49.0003
```


```
NLPID: 0xCC
```

```
Hostname: R5
```

```
IP Address: 10.1.100.5
```

```
Metric: 10 IP 10.1.4.0 255.255.255.0
```

Show ISIS Topology



Shows the List of all L1 routers in the area and L2 routers along with the metric

R4#show isis topology

IS-IS paths to level-1 routers

| System Id | Metric | Next-Hop | Interface | SNPA |
|-----------|--------|----------|-----------|----------------|
| R4 | -- | | | |
| R5 | 10 | R5 | Et1/0 | aabb.cc00.0501 |

IS-IS paths to level-2 routers

| System Id | Metric | Next-Hop | Interface | SNPA |
|-----------|--------|----------|-----------|----------------|
| R2 | 20 | R3 | Et0/0 | aabb.cc00.0300 |
| R3 | 10 | R3 | Et0/0 | aabb.cc00.0300 |
| R4 | -- | | | |
| R5 | 10 | R5 | Et1/0 | aabb.cc00.0501 |

Solution to Scenario 6 : R5 does not have routes to Networks behind R1

```
R5#show isis topology level-2 R2
```

```
Translating "R2"
```

```
IS-IS level-2 path to R2
```

| System Id | Metric | Next-Hop | Interface | SNPA |
|-----------|--------|----------|-----------|----------------|
| R2 | 20 | R4 | E1/0 | aabb.cc00.0401 |



```
R5#sh isis topology R1.00-00
```

```
R5#
```

```
R5#sh isis topology R1
```

```
Translating "R1"
```

```
R5#
```

The node owning this LSP R1 must be isolated from the calculating node

Upon further tracking the Database and topology hop by hop, we could trace that the problem is between R1 and R2

MPLS

Workshop




ISIS Fundamentals and Troubleshooting

Agenda

- Overview
- IS-IS Hierarchical Areas
- Addressing
- MTU and Hello Padding
- Attach-bit and Route Leaking
- LSP Flooding and Convergence
- **SPF and Network Stability**
- Route Redistribution
- Narrow and Wide Metrics

Scenario 7 : Frequent SPF runs and Network Instability





Start with 'Show isis spf-log'

Gives two key information :

- How often are SPF's run ?
 - frequent SPF's can indicate a problem in the network
 - in a stable network SPF should only run periodically
- Who triggered the SPF ?
 - see LSPid of first trigger LSP
 - helps find the source of the problem

Start with 'Show isis spf-log'

```
R4#show isis spf-log
```

```
level 1 SPF log
```

| When | Duration | Nodes | Count | First trigger LSP | Triggers |
|----------|----------|-------|-------|-------------------|-------------------|
| 00:48:48 | 0 | 2 | 1 | | PERIODIC |
| 00:01:39 | 0 | 1 | 2 | R4.00-00 | DELADJ TLVCONTENT |
| 00:01:34 | 0 | 2 | 2 | R4.00-00 | NEWADJ TLVCONTENT |
| 00:01:24 | 0 | 2 | 2 | R5.00-00 | LSPHEADER |

LSP ID of first LSP causing SPF
older IOS shows last trigger LSP

What in LSP triggers ?



Check 'Show isis lsp-log'

- Not the same as “Show isis **spf**-log”
- Gives two key information from LSP
 - How often do we generate new LSP ?
 - Why did it generate a new LSP ?
- One router can potentially influence the whole network
- Flapping adjacency is shown by interface involved

LSP-Log

R4#show isis lsp-log

Level 1 LSP log


| When | Count | Interface | Triggers |
|----------|-------|-------------|--------------|
| 20:12:40 | 1 | Ethernet1/0 | DELADJ |
| 20:12:35 | 1 | Ethernet1/0 | NEWADJ |
| 20:10:18 | | | ATTACHFLAG |
| 19:59:48 | 1 | Loopback0 | CONFIG |
| 19:54:05 | | Loopback0 | IPUP |
| 19:53:53 | 1 | Loopback0 | CONFIG |
| 00:49:11 | 2 | Ethernet1/0 | CLEAR DELADJ |
| 00:49:06 | 1 | Ethernet1/0 | NEWADJ |

event that triggered the LSP to be flooded

up to 20 occurrences are kept

Interface that causes the LSP generation

Debug ISIS spf-triggers



R5#debug isis spf-triggers

IS-IS SPF triggering events debugging is on

Apr 9 15:19:31.179:ISIS-Spf: L1 SPF needed, periodic SPF from 0xA8508
periodic SPF in L1


Apr 9 15:19:31.179:ISIS-Spf: L2 SPF needed, periodic SPF from 0xA8508
periodic SPF in L2

Apr 9 15:11:18.551:ISIS-Spf: L1 SPF needed, L2 attach changed
from 0xA9ED1C
R4 loses L2 with R3

MPLS
Workshop

Solution to Scenario 7 : Frequent SPF runs and Network Instability

Flapping interface on R5 triggering frequent SPFs



Show commands and debugs discussed will help locate the offending LSP in large networks

MPLS
Workshop




ISIS Fundamentals and Troubleshooting

Agenda

- Overview
- IS-IS Hierarchical Areas
- Addressing
- MTU and Hello Padding
- Attach-bit and Route Leaking
- LSP Flooding and Convergence
- SPF and Network Stability
- **Route Redistribution**
- Narrow and Wide Metrics

Scenario 8 : Redistribution to / from IS-IS failing





Gotchas in ISIS Redistribution

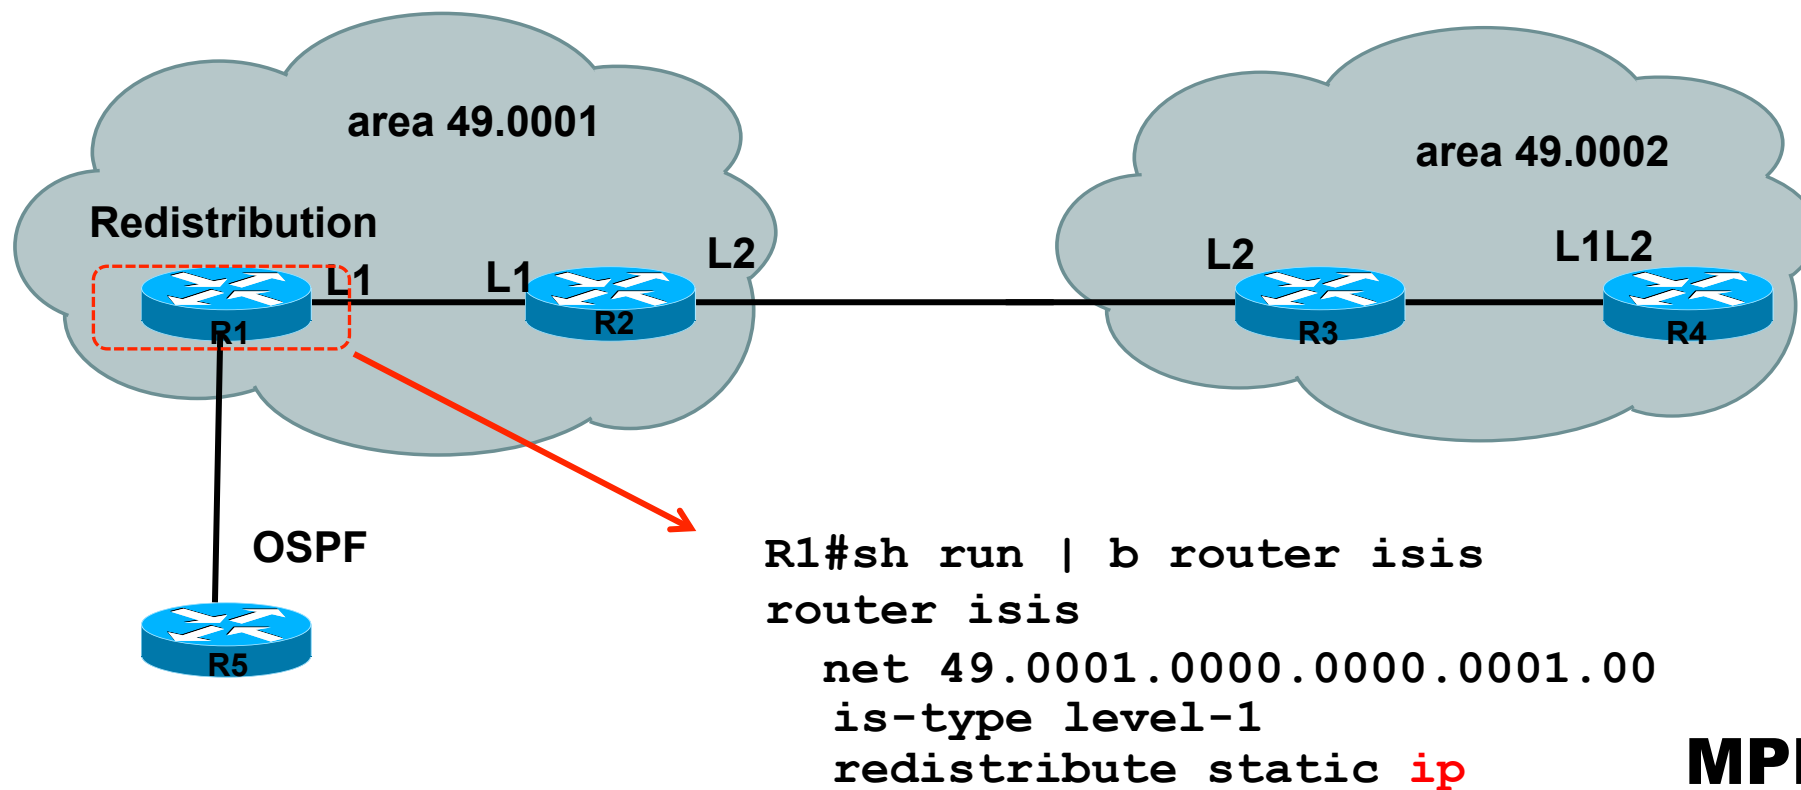
- Default redistribution is for **only level-2** ISIS routes
 - “**redistribute <protocol> level-1**” needs to be specified for L1 routes to get redistributed into any protocol
- “**Redistribute static ip**”
 - IP** keyword needs to be explicitly mentioned, otherwise redistribution of IP prefixes from static into ISIS will fail
- Just like OSPF, ‘subnets’ needs to be added for classless networks to be redistributed

Route Comparison in IS-IS

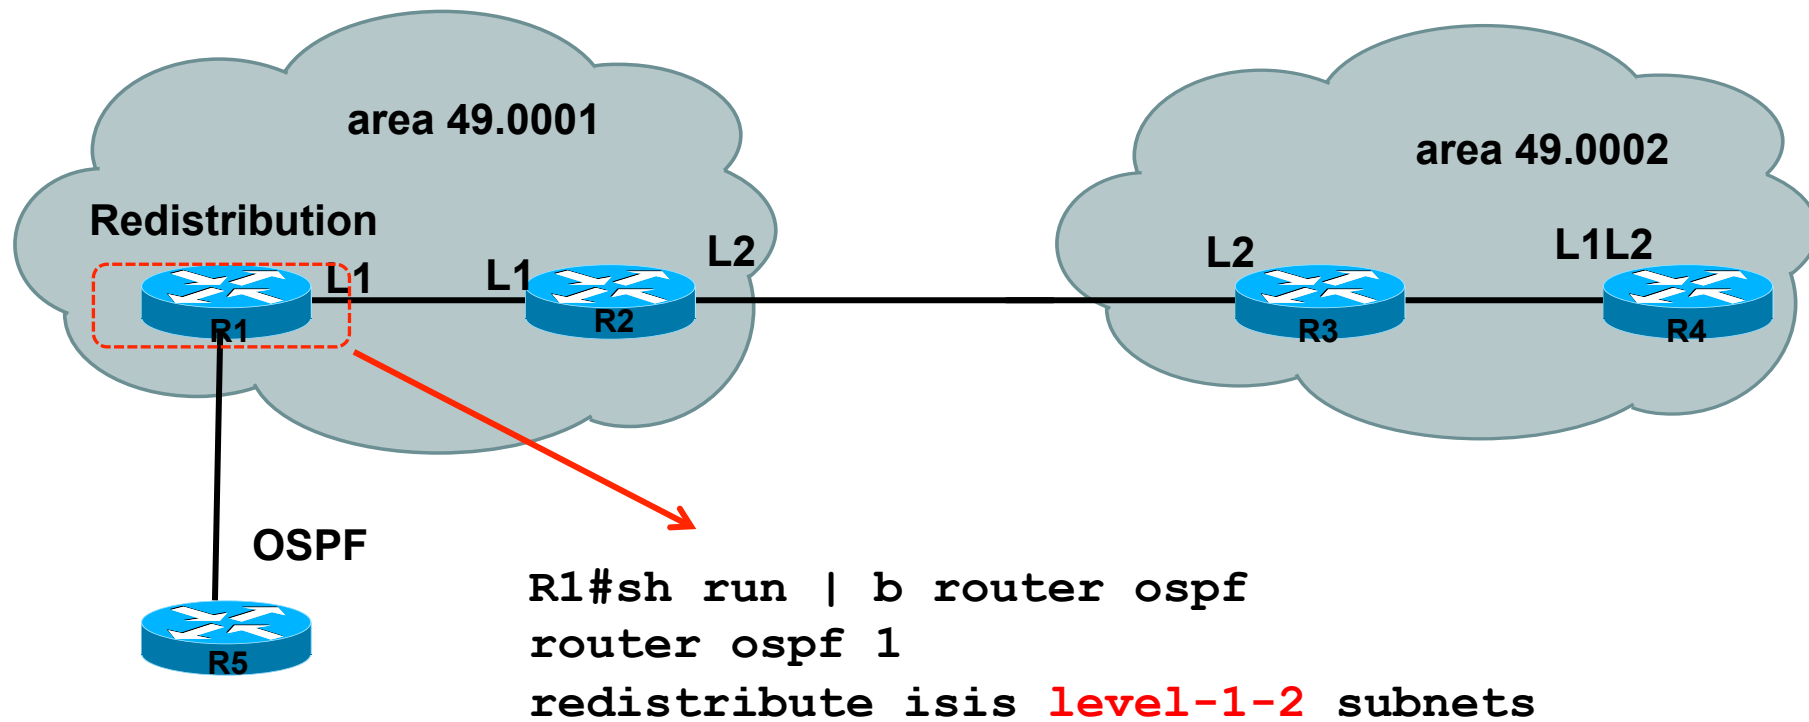


- Between L1 and L2, in case of same routes (same prefix and same mask), Level 1 is preferred over Level 2
- Internal equal to external **route-type** : no difference based on the **route-type** itself, redistributed or internal. Metric-type in the next step decides it.
 - Route type is just based on the TLV used to carry the IP reachability information (TLV 128 versus TLV 130)
- Irrespective of route type, Internal **metric-type** is preferred over external **metric-type**. Metric-type can be set during the redistribution
- Redistributed routes into ISIS are **L1 and Internal** metric-type by default, They compete with regular IS-IS routes right away

Solution to Scenario 8 : Redistribution to / from ISIS failing



Solution to Scenario 8 : Redistribution to / from ISIS failing






ISIS Fundamentals and Troubleshooting

Agenda

- Overview
- IS-IS Hierarchical Areas
- Addressing
- MTU and Hello Padding
- Attach-bit and Route Leaking
- LSP Flooding and Convergence
- SPF and Network Stability
- Route Redistribution
- **Narrow and Wide Metrics**

Scenario 9 :

Database populated, no routes installed on R2



```
R2#sh ip route isis
```

```
R2#
```

```
R2#sh isis database level-1 R1.00.00 detail
```

```
IS-IS Level-1 LSP R1.00-00
```

| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|----------|-------------|--------------|--------------|----------|
| R1.00-00 | 0x00000009 | 0x489F | 515 | 0/0/0 |

```
Metric: 10
```

```
IP 99.99.99.99/32
```

```
Metric: 10
```

```
IS-Extended R1.01
```



TLVs

- TLV = Type Length Value
- Flexible way of carrying information
- New stuff goes into TLVs
- A router that does not recognize a TLV, ignores it

ISIS Packet



| | | | |
|--|---|---|----------|
| Intradomain Routing Protocol Discriminator | | | |
| Length Indicator | | | |
| Version/Protocol ID Extension | | | |
| ID Length | | | |
| R | R | R | PDU type |
| Version | | | |
| Reserved | | | |
| Maximum Area Addresses | | | |
| Additional Header Fields | | | |
| TLV Fields | | | |

TLV format make IS-IS Flexible
in terms of tuning and easily
extendable to new features like
MPLS-TE, IPv6

variable



Narrow and Wide IS-IS Metrics

- Narrow Metrics

- Maximum **LINK_METRIC** per interface is **63** (only 6 bits)
- Maximum **PATH_METRIC** is **1023**

- Wide Metrics

- Extended IS Reachability TLV and Extended IP Reachability TLV introduced wide metrics
- Max **LINK_METRIC** is **16777215** ($2^{24} - 1$)
- Max **PATH_METRIC** is **4261412864** ($2^{32} - 2^{25}$)
- Needed for MPLS Traffic Engineering to work



Mismatch in Metric Style

- Turning on '**metric-style wide**' only on a few routers will break connectivity, as routers running narrow metrics will not understand the newer TLVs
- “**metric-style transition**” will help during Migration - advertise and accept both old and new TLVs

We can use different flavors of transition command, depending on the scenario :

- “**metric-style narrow transition**” - To advertise only old-style TLVs but accept both old and new
- “**metric-style wide transition**” - To advertise only new-style TLVs and accept both

Detecting Mismatch in Metric Style

R1#sh clns protocol

```
IS-IS Router:
Redistribute:
    static (on by default)
Distance for L2 CLNS routes:110
RRR level: none
```

```
Generate narrow metrics: none
Accept narrow metrics: none
Generate wide metrics:level-1-2
Accept wide metrics :level-1-2
```

R2#sh clns protocol

```
IS-IS Router:
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
```




Detecting Mismatch in Metric Style

R2#debug isis update-packets

```
Apr  9 01:46:21.259: ISIS-Upd: Rec L2 LSP 0000.0000.0001.01-00
Apr  9 01:46:21.259: ISIS-Upd: from SNPA 0025.4531.0980
Apr  9 01:46:21.259: ISIS-Upd: LSP newer than database copy
Apr  9 01:46:21.259: ISIS-Upd: TLV code mismatch (22, 2)
Apr  9 01:46:21.259: ISIS-Upd: TID 0 full SPF required
```

Solution to Scenario 9 :

Database populated, no routes installed on R2



```
R2#sh isis database level-1 R1.00.00 detail
LSPID    LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
R1.00-00  0x00000009      0x489F        515           0/0/0
Area Address: 49.0001
Metric: 10      IP 172.16.1.0/24
Metric: 10      IP 99.99.99.99/32
Metric: 10      IS-Extended R1.01
```


Review : ISIS vs OSPF

| | ISIS | OSPF |
|----------------------|--|--|
| Design | 2-level hierarchy Originally: Only totally stub areas route leaking makes areas non-stub | 2-level hierarchy Multiple types or areas |
| Encapsulation | Runs directly over layer 2 | On top of IP (can be remotely attacked) |
| Flooding | Reliable on p2p links 1 LSP per router Flexibility via TLVs Unrecognized TLVs are flooded Requesting info and acks (PSNP) DIS | Reliable Many different types (11) of LSAs Fixed LSA formats Unrecognized LSA are not flooded Requesting info and acks (LS Request/Update) DR and BDR |
| Scalability | Thousands of prefixes | Thousands of prefixes |

Review : ISIS vs OSPF

| | ISIS | OSPF |
|-----------------|---|---|
| Network types | P2P Broadcast | P2P Broadcast Non-broadcast Point-to-multipoint |
| Aging | Periodic flooding Aging: counts down Remaining lifetime is configurable Cannot disable aging | Periodic flooding Aging: counts up Maxage is not configurable DoNotAge (DNA) bit cancels aging out |
| MPLS TE support | Yes | Yes |
| IPv6 | Integrated | Seperate and new protocol: OSPFv3 |
| Authentication | Yes | Yes |

Further Reading




LTRARC-2002

© 2015 Cisco and/or its affiliates. All rights reserved.

Cisco Public

12
1



MPLS
Workshop



APPENDIX

MPLS
Workshop



IS-IS Common Recommendations

- Unless required, CLNS should be disabled with “**no clns routing**” for IP-only networks
- The overload bit should be set with “**set-overload-bit**” to protect against traffic being “black holed” on initial router bootup. Generally the recommended timeframe is 180 seconds (3 minutes)
- Ensure the router is set to ignore LSP errors with “**ignore-lsp-errors**”, like checksum errors to avoid overload on originating router due to perpetual purge and regeneration – enabled by default
- “**Log-adjacency-changes**” should be enabled on all platforms across the network to ensure easier troubleshooting
- Ensure “**hostname dynamic**” command is enabled to create system-ID-to-router-name mapping for easier troubleshooting – enabled by default



ISIS Master show command List

show commands

show clns
show clns interface
show clns neighbor
show clns neighbor detail
show clns protocol
show clns route
show isis database
show isis database detail
show isis hostname
show isis spf-log
show isis topology
show isis route
show clns traffic
show clns cache

Hidden show commands

show isis timers
show isis private
show isis database private
show isis spf-log detail
show isis tree
show isis lsp-log



Thank you.

MPLS
Workshop



CISCO