
TOMORROW
starts here.
CISCO
Cisco live!

ISIS Deployment in Modern Networks

BRKRST-2338

Mani Ganesan – CCIE R&S / SP #27200
Consulting Systems Engineer
@mani_cisco

Cisco live!

Agenda

- ISIS Overview
  - CLNS, L1/L2 Routing, Best Practices
- ISIS for IPv6
  - Single Topology, Multi-Topology
- ISIS in the Backbone
  - Fast Convergence Features
- ISIS at the Edge
  - BGP and MPLS Considerations
- ISIS at the Access / Aggregation
  - Route Leaking, Traffic Engineering and IP FRR

Reference only slide

Cisco live!

What is IS-IS ?

Intermediate System-to-Intermediate System (IS-IS) Overview

- IS-IS is a link-state routing protocol;
  - Commonly used in Service Providers and large Enterprise networks.
  - Offer Fast convergence
  - Excellent scalability
  - Flexibility in terms of tuning
- Easily extensible with Type/Length/Value (TLV) extensions;
  - IPv6 Address Family support (RFC 2308)
  - Multi-Topology support (RFC 5120)
  - MPLS Traffic Engineering (RFC 3316)

What is IS-IS ?

CLNS Encapsulation of IS-IS

- IS-IS is a Layer 2 protocol and is not encapsulated in IP
- Logical Link Control (LLC) 802.3 Data-link header for IS-IS uses :
  - DSAP (Destination Service Access Point) set to 0xFE
  - SSAP (Source Service Access Point) set to 0xFE
- IS-IS Fixed header
- IS-IS Data encoded as Type-Length-Value (TLV)

What is IS-IS ?

IS-IS Addressing

- Each IS-IS router is identified with a Network Entity Title (NET)
- ISPs commonly choose addresses as follows:
  - First 8 bits – pick a number (49 used in these examples)
  - Next 16 bits – area
  - Next 48 bits – router loopback address
  - Final 8 bits – zero
- Example:
  - NET: 49.0001.1921.6800.1001.00
  - Router:192.168.1.1(loopback) in Area1

IOS Example:

```
!
interface Loopback0
  ip address 192.168.1.1 255.255.255.255
!
interface Ethernet0
  ip address 192.168.12.1 255.255.255.0
  ip router isis
!
router isis
  passive-interface Loopback0
  net 49.0001.1920.1680.1001.00
```

ISIS vs OSPF

* draft-bhatia-manral-diff-isis-ospf

Notable Similarities and Differences

IS-IS and OSPF are both link state protocols, there are similarities and differences

- Similarities:
  - Link-state representation, aging, and metrics
  - Use of Link-state databases and SPF algorithms
  - Update, routing decisions, and flooding processes similar
- Differences:
  - IS-IS organizes domain into two layers; OSPF designates backbone area (area 0)
  - IS-IS peering is more flexible than OSPF (hello time, dead intervals, and subnet mask need not match)
  - IS-IS selects single DIS which may be preempted; OSPF elects a DR/BDR which cannot be preempted,
  - IS-IS does not support NBMA, point-to-multipoint, or virtual links (it rides L2 directly)

ISIS vs OSPF

Terminology

- OSPF
  - Host
  - Router
  - Link
  - Packet
  - Designated router (DR)
  - Backup router (BDR)
  - Links State Advertisement (LSA)
  - Hello Packet
  - Database Description (DBD)

- ISIS
  - End System (ES)
  - Intermediate System (IS)
  - Circuit
  - Protocol Data Unit (PDU)
  - Designated IS (DIS)
  - N/A (no DBIS is used)
  - Link State PDU (LSP)
  - IIH PDU
  - Complete Sequence Number PDU (CSNP)

ISIS vs OSPF

OSPF Areas - Example

- OSPF
  - Area
  - Backbone Area (area 0)
  - Non-backbone area
  - Area Border Router (ABR)
  - Autonomous System Boundary Router (ASBR)

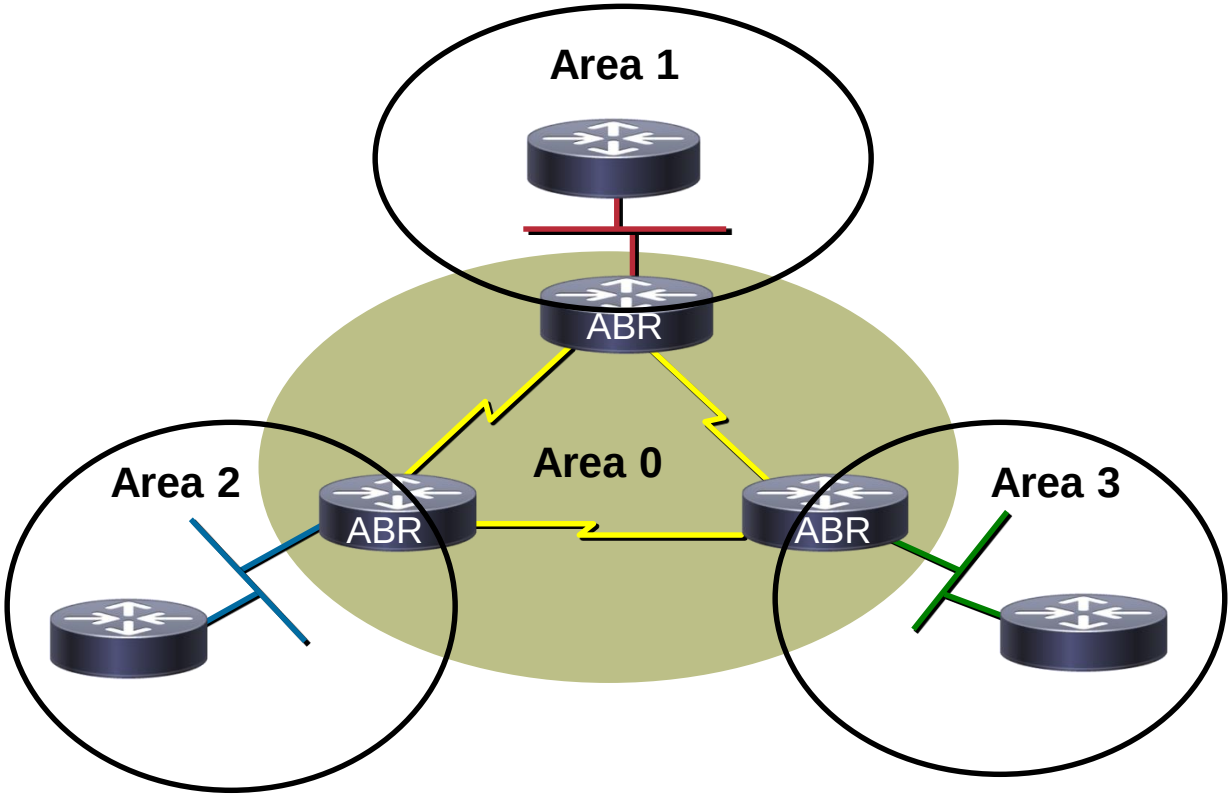

ISIS vs OSPF

ISIS Areas - Example

- ISIS
  - Sub domain (area)
  - Level-2 Sub domain (backbone)
  - Level-1 area
  - Level-1-2 router (L1-L2)
  - AS boundary can be any router (IS)
- IS-IS does not have back-bone “area”
  - A backbone is a contiguous collection of Level-2 routers

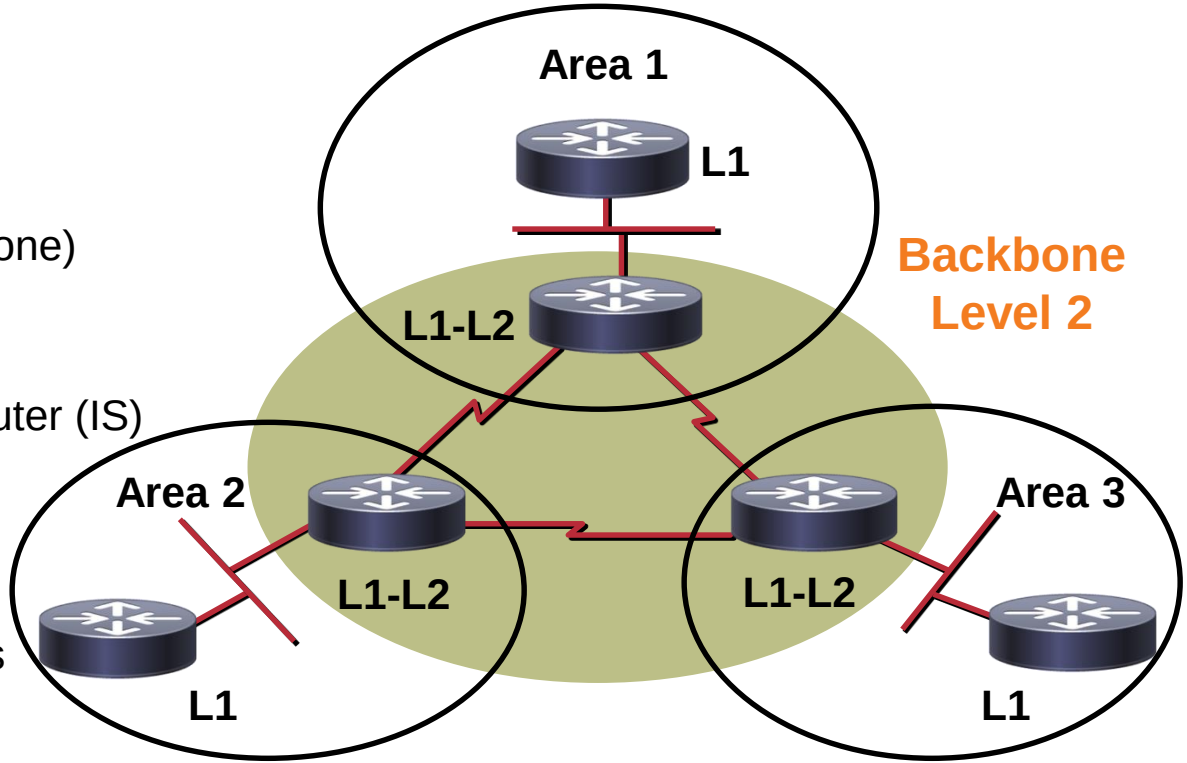

Hierarchy Levels

- IS-IS presently has a two-layer hierarchy
  - The backbone (level 2)
  - Non-backbone areas (level 1)
- An IS (router) can be either:
  - Level 1 router (used for intra-area routing)
  - Level 2 router (used for inter-area routing)
  - Level 1–2 router (intra and inter-area routing) (by default Cisco routers are L1-L2)

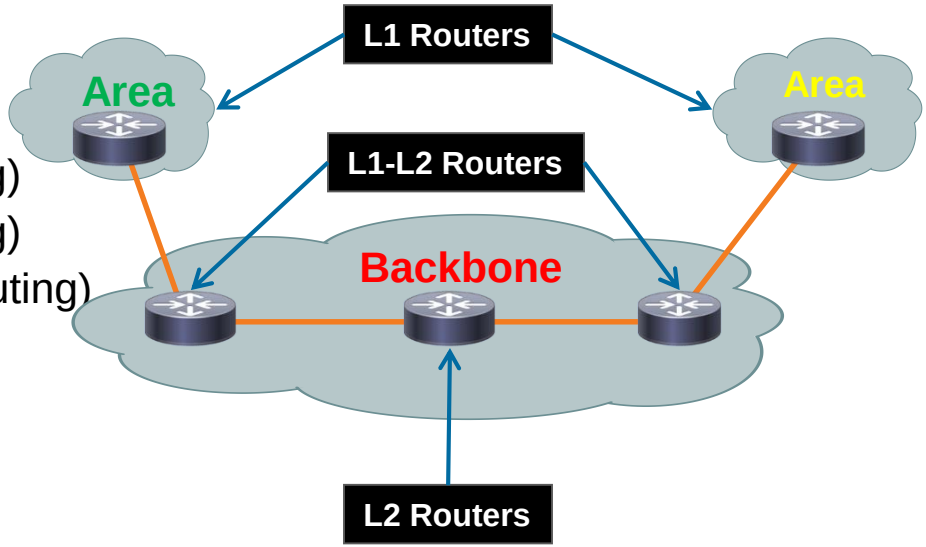

Hierarchy Levels

Level 1 Routers

- Level 1-only routers
  - Can only form adjacencies with Level 1 routers with-in the same area
  - Link State Data Base (LSDB) only carries intra-area information

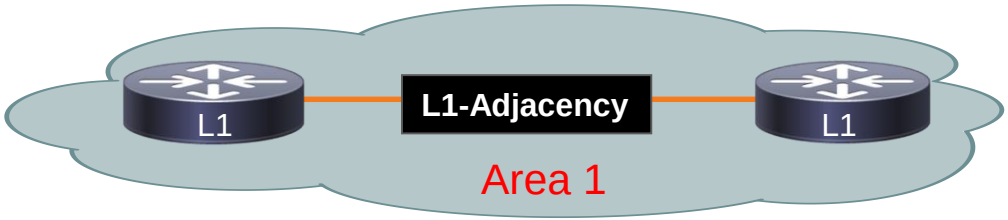

Hierarchy Levels

Level 2 Routers

- Level-2-only routers
  - Exchange information about the L2 area
  - Can form adjacencies in multiple areas

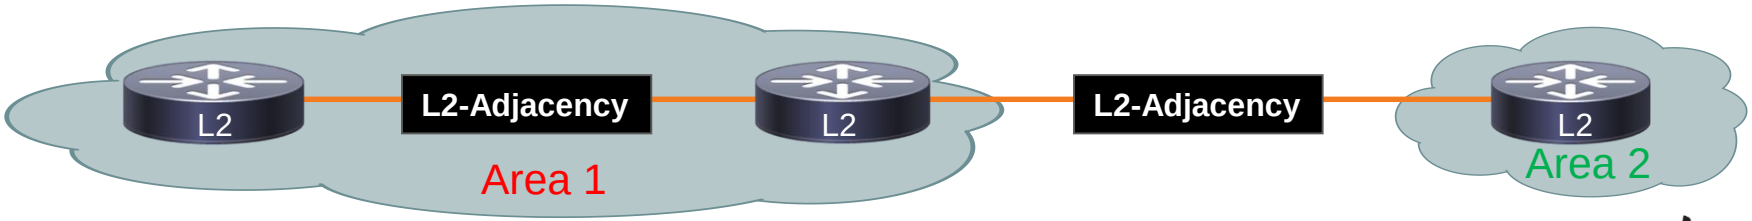

Hierarchy Levels

Level1-2 Routers

- Level 2 routers may also perform Level-1 routing (L1-L2 routers)
  - L1-L2 routers may have neighbors in any area
  - Have two separate LSDBs: Level-1 LSDB & Level-2 LSDB
- Level 1-2 routers carry other L1 area information;
  - How to reach other L1 areas via the L2 topology
  - Level 1 routers look at the Attached-bit (ATT-bit) to find the closest Level 1–2 router
  - Installs a default route to the closest Level 1–2 router in the area

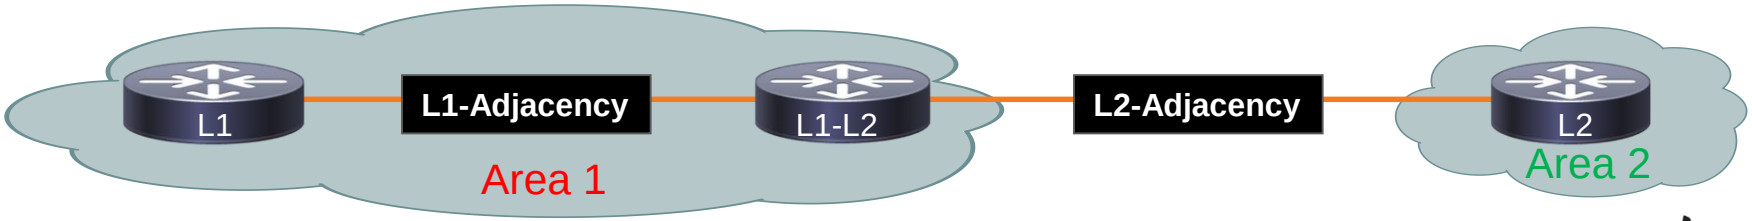

Hierarchy Levels

Level 1, Level 2, and Level 1–2 Routers

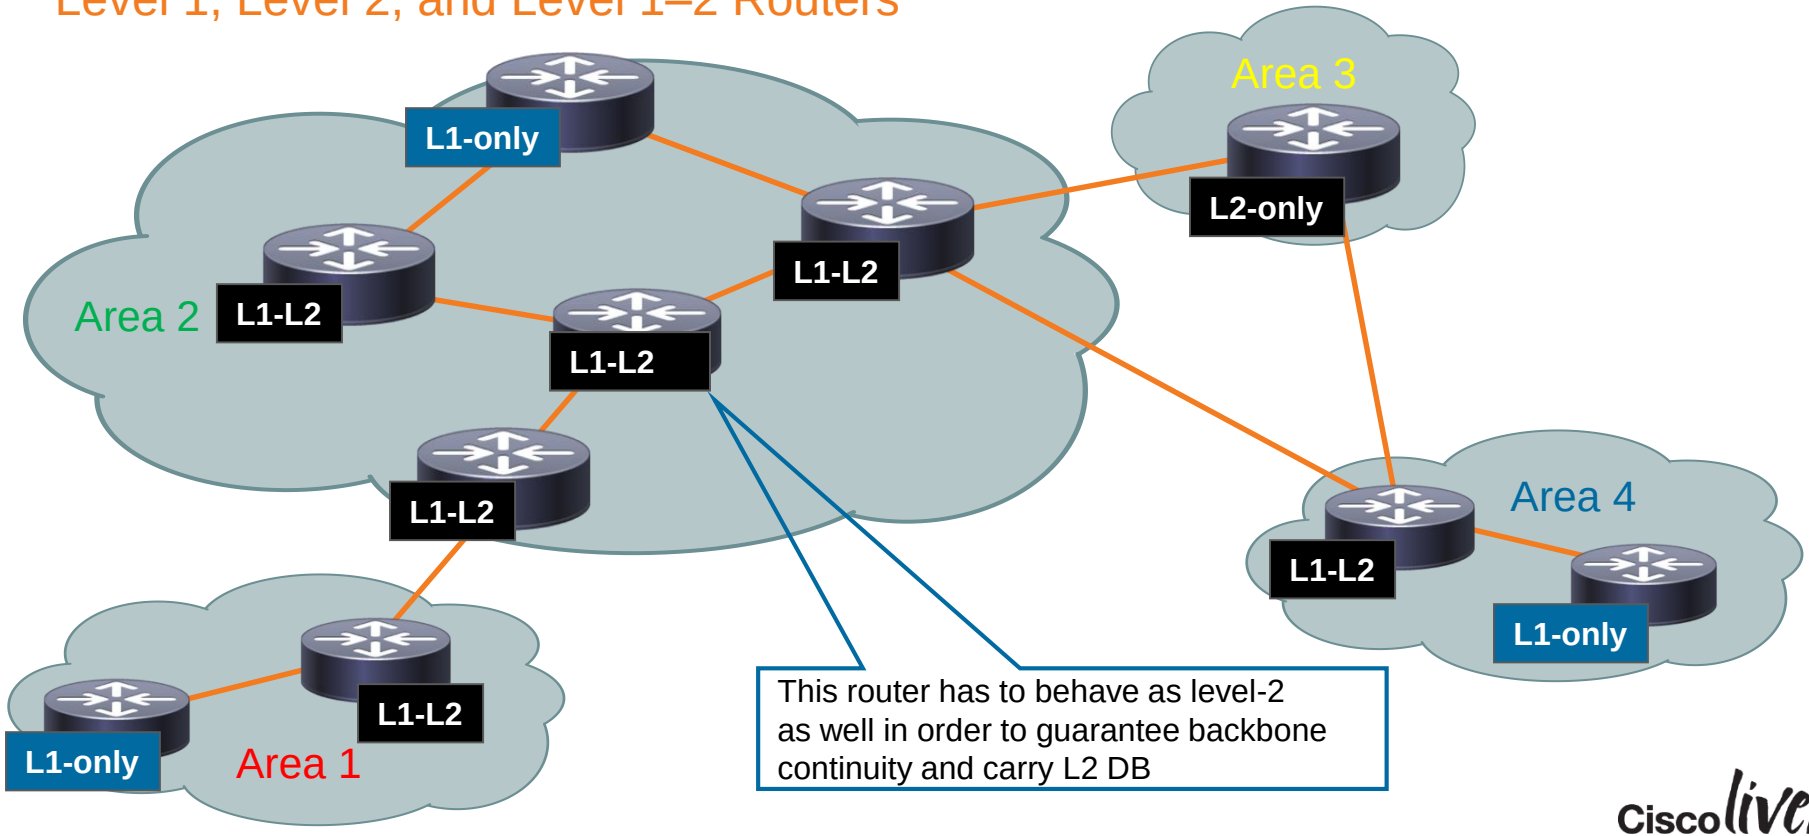

ISIS Overview
- Best practices

Cisco live!

Setting IS-IS Metric

- ISIS interface cost is not dynamic and there is no auto-cost reference, the default metric for all interfaces is 10 for both L1 and L2

- Manually configure Metric across the network with "isis metric" interface command according to overall routing strategy
  - Compare with OSPF which set cost according to link bandwidth

- If a link, such as one that is used for traffic engineering, should not be included in the SPF calculation, enter the isis metric command with the maximum keyword.

Increase IS-IS Default Metric

- Keeping the default metric as 10 across the network is not optimal, if configured value on any preferred interface is “accidentally” removed - a low priority interface could end up taking full load by mistake
- Configure a “very large” value as default across the network - metric 100000

- Summary address cost :
  - The best available cost from the more specific routes (plus cost to reach neighbor of the best specific).
  - Adjust the cost of the best specific route to control if summarizing at different points.

IS-IS MTU Mismatch detection

Disable Hello padding

- Disable IS-IS Hello [IIH] padding
  - On high speed links, it may strain huge buffers
  - On low speed links, it waste bandwidth
  - May affect time sensitive applications, e.g., voice
- IOS will pad the first 5 IIH's to the full MTU to aid in the discovery of MTU mismatches.

```
router isis
 no hello padding
```

Or

```
interface <>
  no isis hello padding
```

- "Sometimes" option on IOS-XR will use hello padding for adjacency formation only

Agenda

- ISIS Overview
  - CLNS, L1/L2 Routing, Best Practices
- ISIS for IPv6
  - Single Topology, Multi-Topology
- ISIS in the Backbone
  - Fast Convergence Features
- ISIS at the Edge
  - BGP and MPLS Considerations
- ISIS at the Access / Aggregation
  - Route Leaking, Traffic Engineering and IP FRR

ISIS for IPv6

- Single Topology, Multi-Topology

Cisco live!

IS-IS for IPv6

- IPv6 Address Family support (RFC 2308)
- 2 new Tag/Length/Values added to introduce IPv6 routing
  - IPv6 Reachability TLV(0xEC):
    - Equivalent to IP Internal/External Reachability TLV's
  - IPv6 Interface Address TLV(0xE8)
    - For Hello PDUs, must contain the link-local address
    - For LSP, must contain the non-link local address
- IPv6 NLPID (Network Layer Protocol Identifier) (0x8E) is advertised by IPv6 enabled routers

IS-IS for IPv6-Only –Example

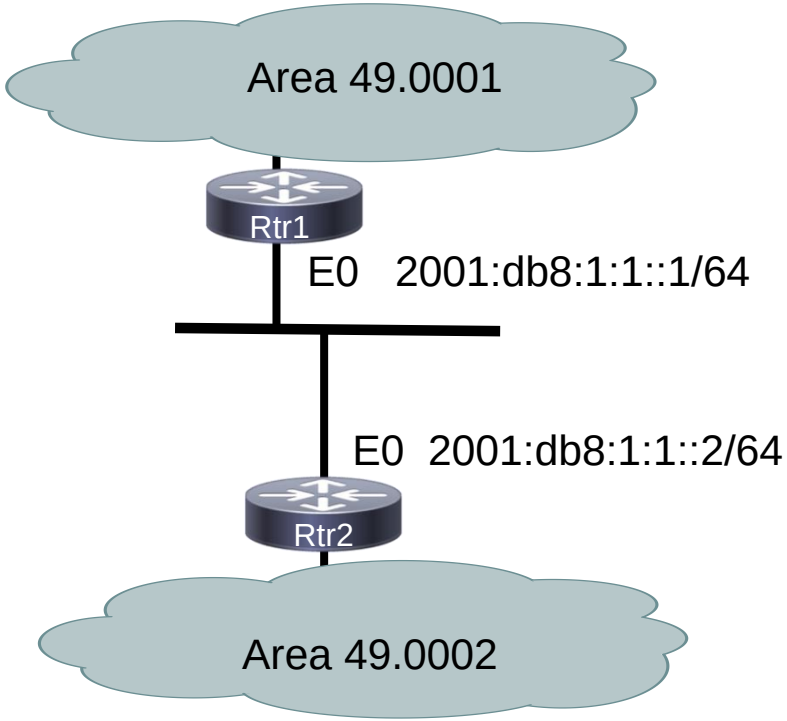

IOS

```
Rtr1#
ipv6 enable
interface ethernet0
 ipv6 address 2001:db8:1:1::1/64
 ipv6 router isis
 isis circuit-type level-2-only
!
router isis
  net 49.0001.1921.6801.0001.00
  address-family ipv6
    redistribute static
  exit-address-family
```

ASR9K

```
Rtr2#
interface ethernet0
 ipv6 address 2001:db8:1:1::2/64
 ipv6 enable
!
router isis
  net 49.0001.1921.6802.0001.00
  address-family ipv6 unicast
    single-topology
    redistribute static
  exit-address-family
  interface fastethernet0/0
    circuit-type level-2-only
    address-family ipv6 unicast
```

IS-IS with dual stack - IOS Example

Dual IPv4/IPv6 configuration

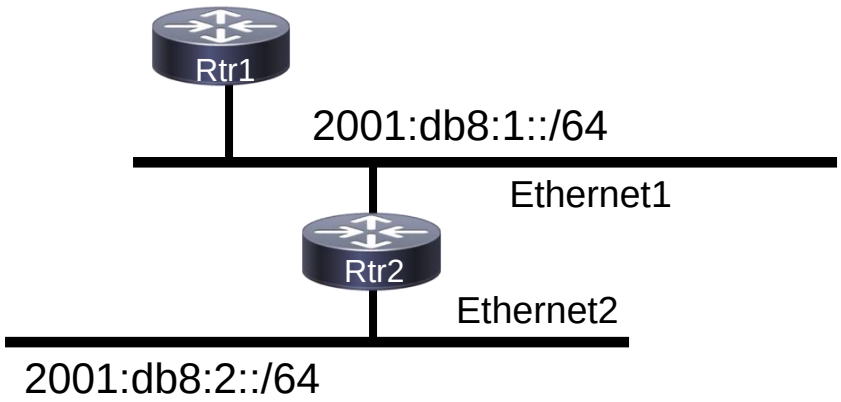

Redistributing both IPv6 static routes and IPv4 static routes.

```
Rtr1#
interface ethernet1
 ip address 10.1.1.1 255.255.255.0
 ipv6 address 2001:db8:1::1/64
 ip router isis
 ipv6 router isis

interface ethernet2
 ip address 10.2.1.1 255.255.255.0
 ipv6 address 2001:db8:2::1/64
 ip router isis
 ipv6 router isis

router isis
 net 49.0001.0000.0000.072c.00
 address-family ipv6
  redistribute static
 exit-address-family
redistribute static
```

IS-IS for IPv6

Restrictions with Single Topology

- In Single topology IS-IS for IPv6 uses the same SPF for both IPv4 and IPv6.
  - IPv4 and IPv6 topologies MUST match exactly
  - Cannot run IS-IS IPv6 on some interfaces, IS-IS IPv4 on others.
  - An IS-IS IPv6-only router will not form an adjacency with an IS-IS IPv4/IPv6 router (Exception is over L2-only interface)

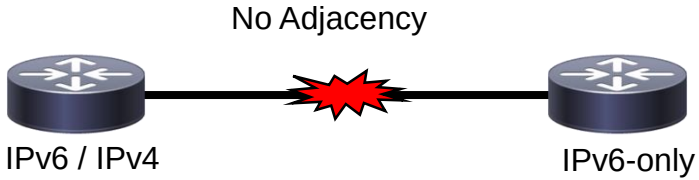

- Cannot join two IPv6 areas via an IPv4-only area
  - L2 adjacencies will form OK
  - IPv6 traffic will black-hole in the IPv4 area.

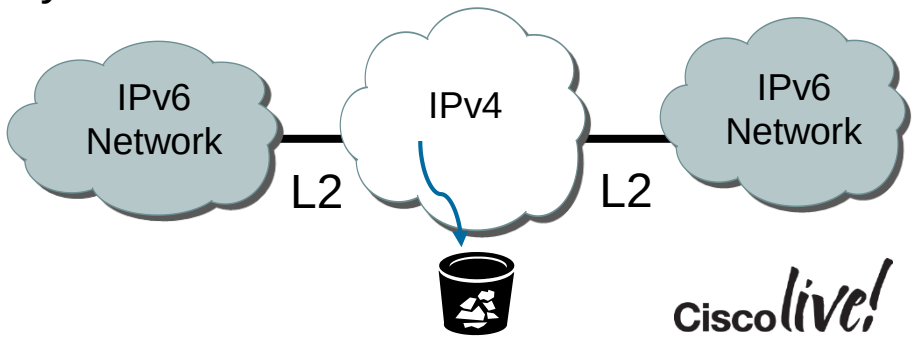

IS-IS for IPv6

Multi-Topology IS-IS extensions

- Multi-Topology IS-IS solves the restrictions of Single topology
  - Two independent topology databases maintained
  - IPv4 uses Multi-Topology ID (MTID) zero(0)
  - New Multi-Topology ID (MTID #2) for IPv6
- Multi-Topology IS-IS has updated packets
  - Hello packets marked with MTID #0 or MTID #2
  - New TLV attributes introduced
  - Each LSP is marked with the corresponding MTID
- Miss-Matched MTID values
  - No effect on broadcast segments, adjacency will form
  - Point-to-point segments, adjacency will not form

```
router isis
 net 49.0001.0000.0000.072c.00
 metric-style wide
 !
 address-family ipv6
  multi-topology
 exit-address-family
```

IS-IS for IPv6

Choosing Single or Multi-Topology IS-IS

- Use Single-Topology (IOS default) for;
  - No planned differences in topology between IPv4 and IPv6
  - Each interface has the same IPv4 and IPv6 router Level

- Use Multi-Topology for;
  - Incremental roll-out of IPv6 on an IPv4 topology
  - If you plan for differences in topology between IPv4 and IPv6

- The optional keyword transition may be used for transitioning existing IS-IS IPv6 single Topology mode to Multi-Topology IS-IS

IS-IS for IPv6

Transition to Multi-Topology IS-IS – Wide Metrics

- Ensure “Wide metric” is enabled
  - Mandatory for Multi-Topology to work
  - When migrating from narrow to wide metrics, care is required
  - Narrow and wide metrics are NOT compatible with each other

- Migration is a two stage process
  - Step 1: make use of the transition keyword

```
router isis
  metric-style transition
```

```
router isis
  metric-style wide
```

  - Step 2: Once the whole network is changed to transition support, the metric style can be changed to wide

Agenda

- ISIS Overview
  - CLNS, L1/L2 Routing, Best Practices
- ISIS for IPv6
  - Single Topology, Multi-Topology
- ISIS in the Backbone
  - Area Design, Fast Convergence Features
- ISIS at the Edge
  - BGP and MPLS Considerations
- ISIS at the Access / Aggregation
  - Route Leaking, Traffic Engineering and IP FRR

ISIS in the Backbone

- Area design

Cisco live!

Area and Scaling

Areas vs. single area

- ISIS supports a large number of routers in a single area
  - More than 400 routers in the backbone is possible
- Starting with L2-only everywhere is a good choice
  - Backbone continuity is ensured from the start
  - Future implementation of level-1 areas will be easier
- Use areas in places where sub-optimal routing is acceptable
  - areas with a single exit point is a better choice from an optimal routing standpoint

Areas and Suboptimal Routing

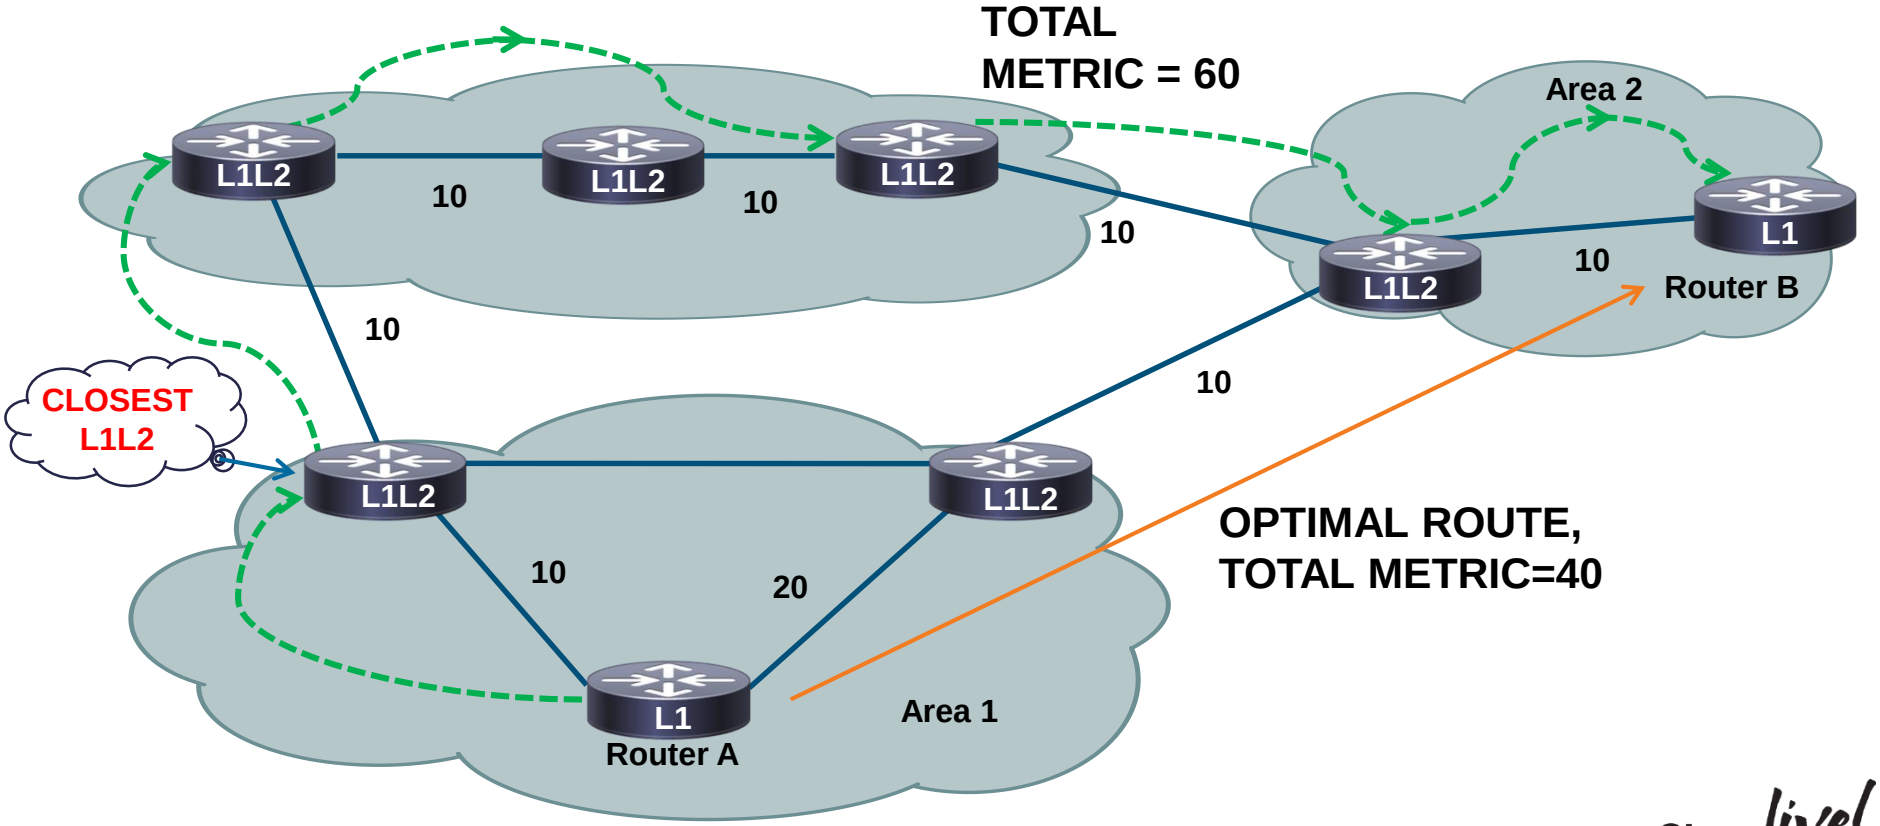

Area Design

L1-Only POP

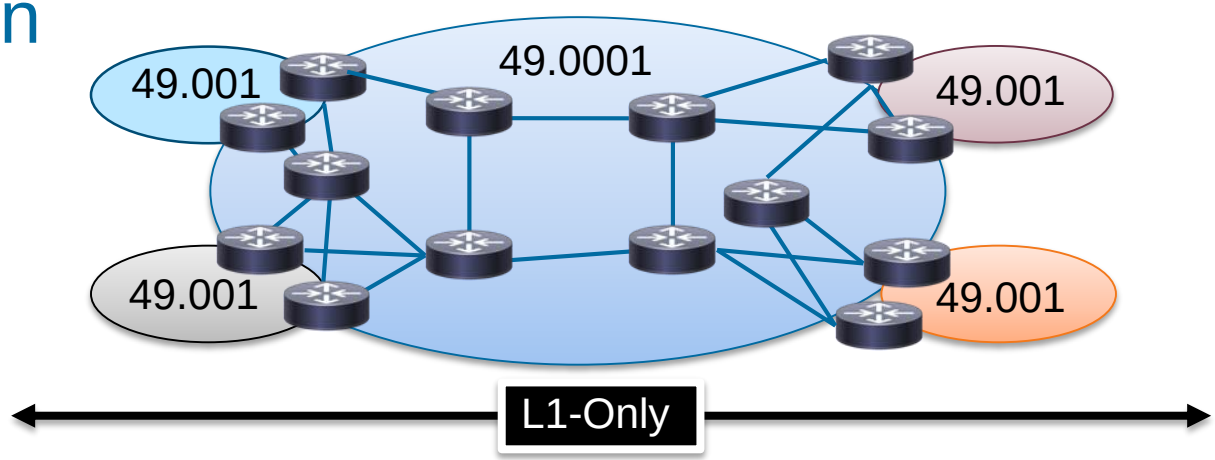

- In this design, all the routers will be running in one area and are all doing L1-only routing
- This design is flat with a single L1-only database running on all the routers
- If you have a change in the topology, the SPF computation will be done in all the routers as they are in the L1-only domain
- SPs picked L1-only to avoid sub-optimal routing problems

Area Design

L2-Only POP

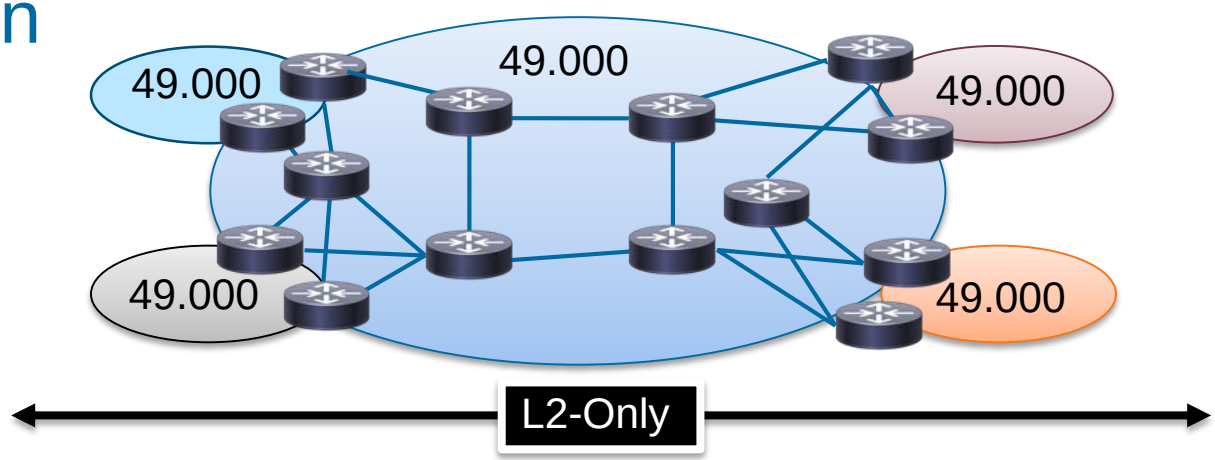

- In this design, all the routers will be running L2-Only in the network
  - With the same Area in all the POPs
- Optimal routing with L2-only database
- Traffic-engineering support with no restrictions, just like L1-only

Area Design

L2-Only POP

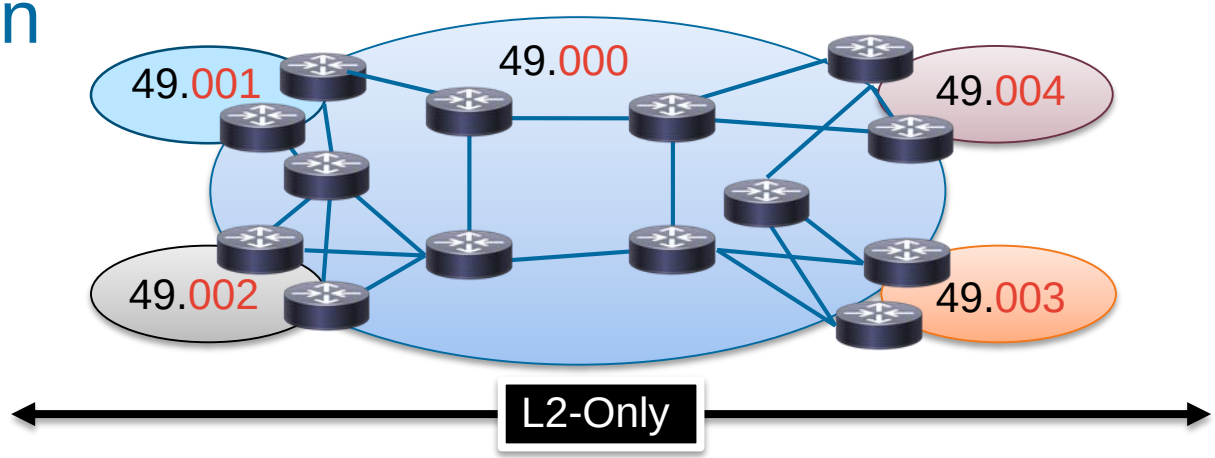

- In this design, all the routers will be running L2-Only in the network
  - With the different Area in all the POPs
  - No summarization and No route-leaking
- All the routers in L2 will share all the LSPs and provides optimal routing (similar to L1-Only POPs)
- As the network grows, easy to bring the L1-only POPs/sub-networks for easy migration

Area Design

L1 in the POP and L2 in the Core

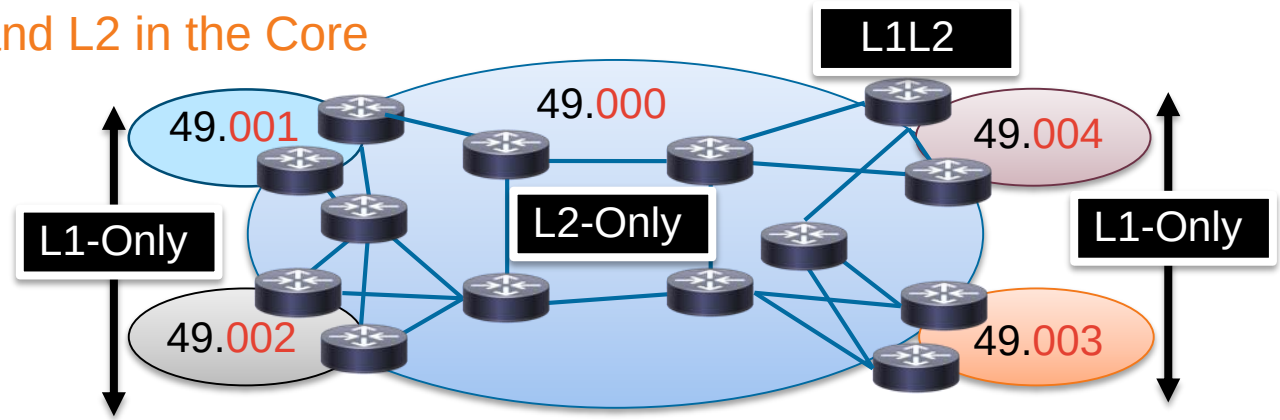

- Within a given local pop—all the routers will be in a separate area
- The L1-L2 routers at the edge of the POPs will be running
  - L1-adj going into the POP
  - L2-adj into the core with the rest of the L1-L2 routers
- The SPF computations will be limited to the respective L1-areas only

Area Design

L1 in the POP and L2 in the Core

- All the L1-routers in a given pop will receive the ATT bit set by the L1L2 router at the edge of the POP
  - L1 routers install a default route based on the ATT bit
- This will cause sub-optimal routing in reaching the prefixes outside the POP by the local routers
- Summarization at the L1L2 boundary
  - potential sub-optimal inter-area routing in certain failure conditions
  - potential black-holing of traffic
  - potential breaking of MPLS LSP among PEs

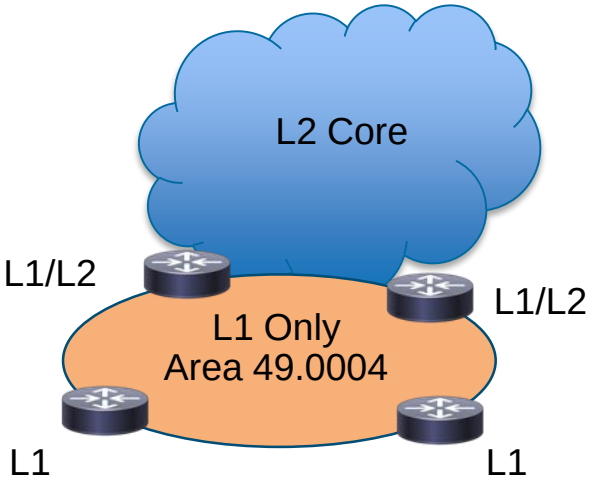

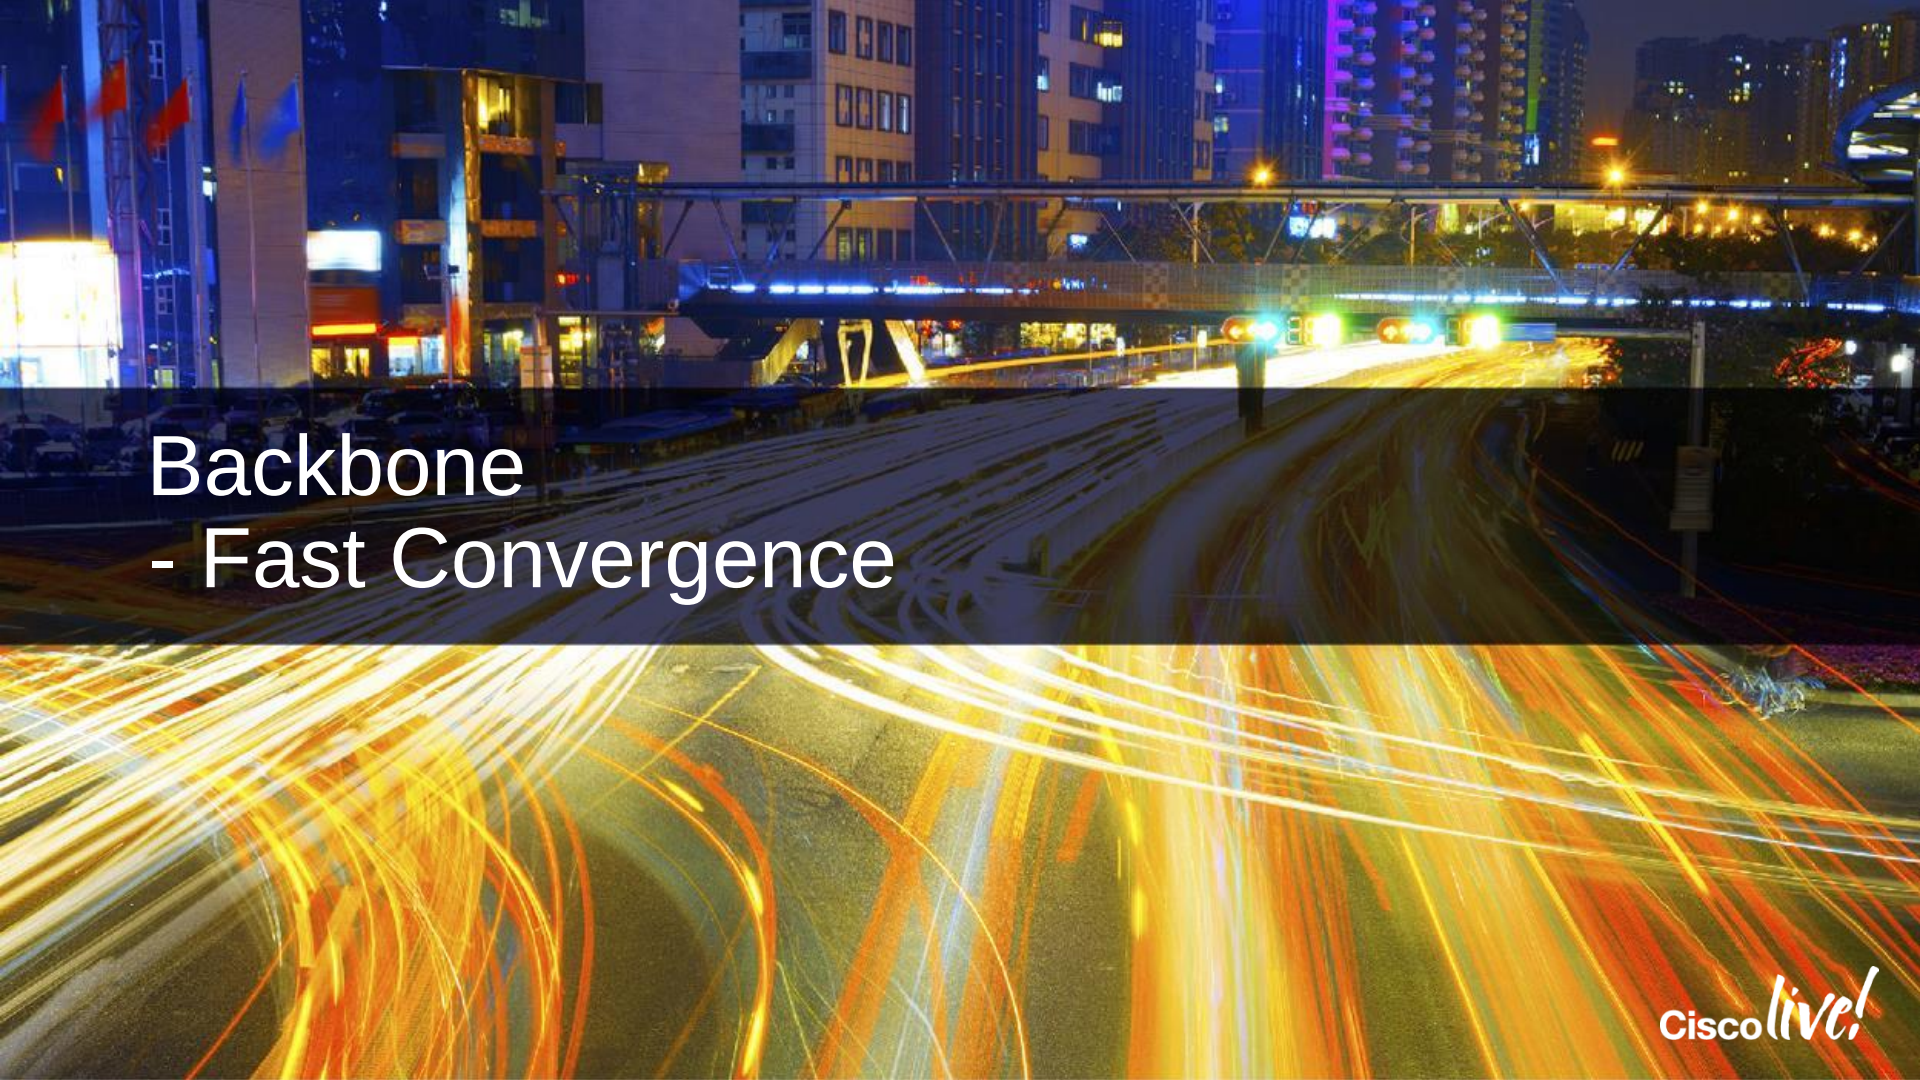

Backbone
- Fast Convergence

Convergence - Overview

- Assume a flow from A to B
- T1: when L dies, the best path is impacted
  - loss of traffic
- T2: when traffic reaches the destination again
- Loss of Connectivity: T2 – T1, called “convergence”

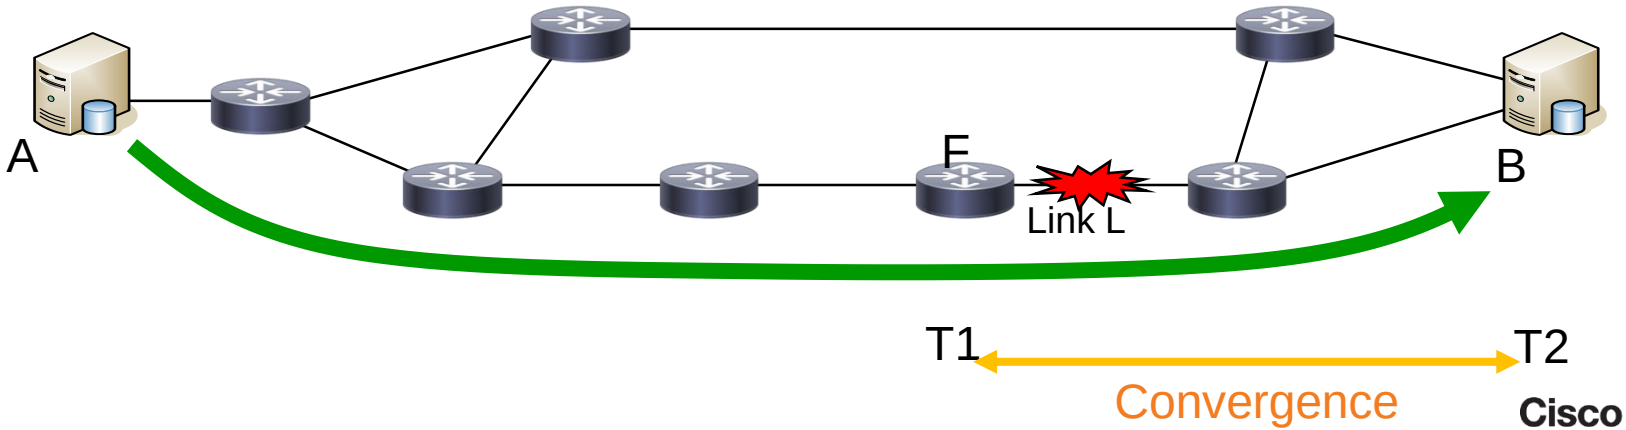

IS-IS Fast Convergence

- Historical IGP convergence ~ O(10-30s)
  - Focus was on stability rather than fast convergence
- Optimizations to link state IGPs enable reduction in convergence to ~ms range with no compromise on network stability or scalability
  - Enables higher availability for all classes of traffic
- If fast reroute techniques are used, traffic restoration may happen well before the network convergence.

IS-IS Fast Convergence

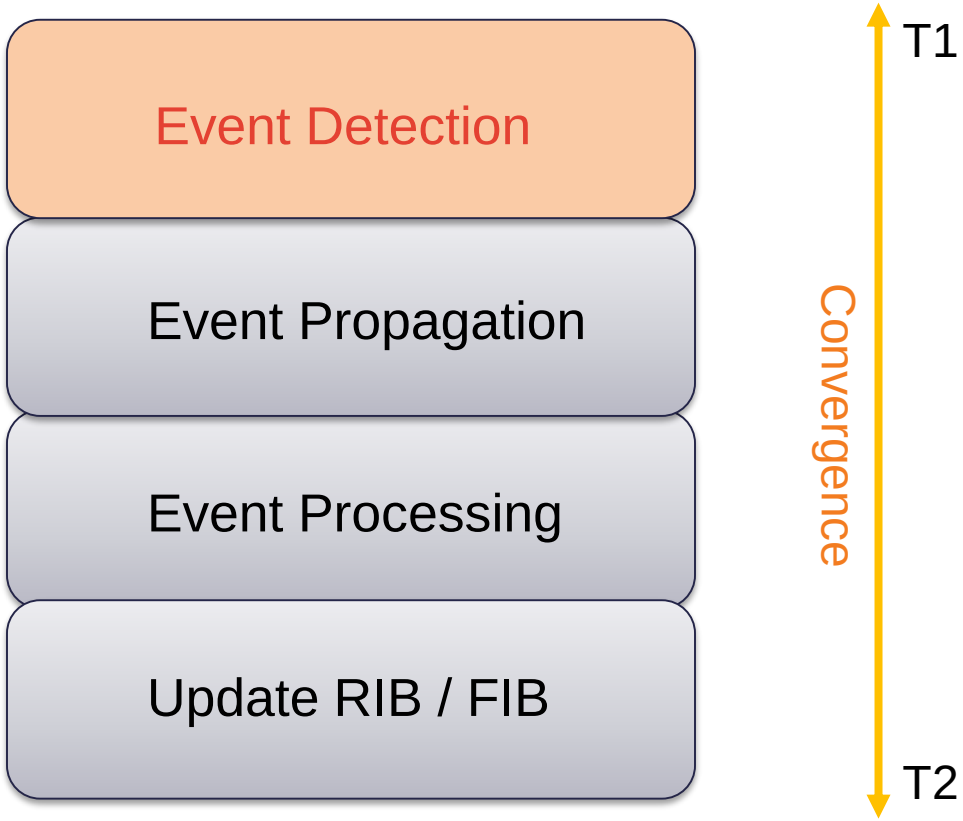

Event Detection

Link Neighbor Failure Detection

- Indirect link failures take time to detect
- With no direct HW notification of link or node loss, convergence times are dependent on Routing Protocol Hellos

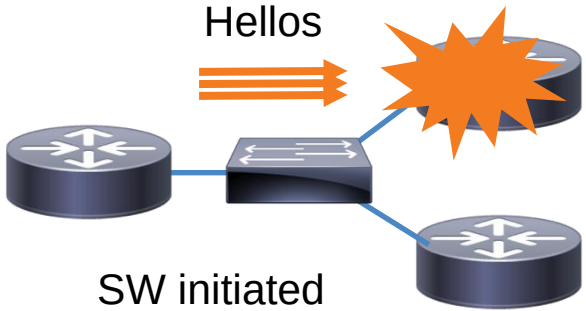

- Hardware detection and recovery is both faster and more deterministic
- Use point-to-point routed links in the Core!

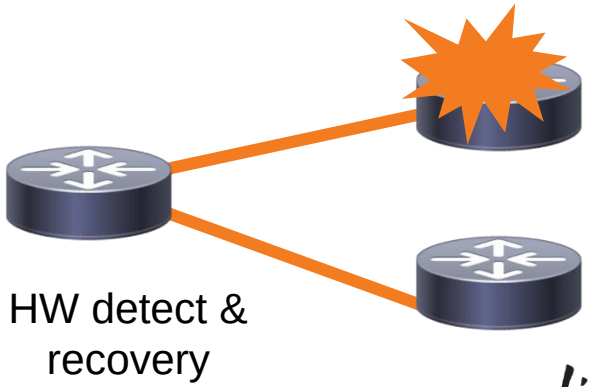

Event Detection

- POS – Best for Convergence
  - Very fast Link failure detection
  - Native anti-flap property of POS
  - down info is signaled very fast
  - up info is confirmed for 10s before relaying to interface

- Other types of Links
  - When physical interface changes state, driver must notify the routing process
  - this should happen in ms range
  - carrier-delay is configurable (Recommendation on IOS : 0 down, 2s Up)
  - IOS-XR and NX-OS have carrier delay of 0 by default

Failure Detection with BFD

- Bidirectional Forwarding Detection (BFD)* provides a lightweight protocol independent mechanism, Improving Indirect Layer 3 Neighbor Failure Detection
  - With BFD running on the interface, a failure of the link would signal IS-IS immediately

```
interface GigabitEthernet 4/1
  bfd interval 100 min_rx 100 multiplier 3
!
router isis
 bfd all-interfaces
```

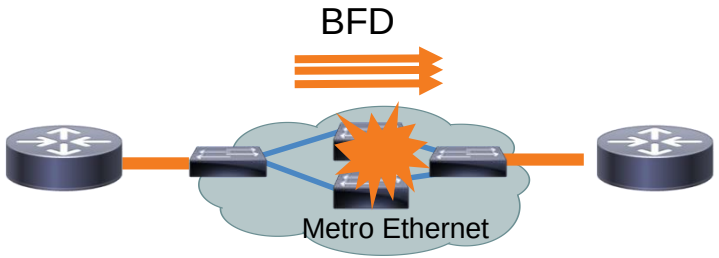

Event Detection

Fast Hellos

- Same Hellos sent more frequently !
  - ~1 second detection
- Process Driven (Scheduler)
- Different Hello per Protocol
  - PIM, LDP, IS-IS, OSPF..
- Handled by Central CPU
  - False positives and load to CPU
- Bandwidth intensive - 50+ Bytes

BFD

- Protocol independent, Even Faster
  - 50ms x 3 = 150ms
- Interrupt Driven like CEF (no waiting)
- Single Hello Type
  - Clients are IS-IS, OSPF..
- Hardware Offloaded on some platforms
  - Nexus, ASR 1k/9k, 7600 ES+
- Light weight ~24 bytes

Enable P2P adjacency over LAN

- When LAN interfaces are used between two routers, we can configure ISIS to behave as p2p
  - Avoid DIS election
  - Avoid CSNP transmissions
- One step less in SPF computation and reduced number of nodes in SPT (no pseudonode)

```
int GigabitEthernet 4/1
  isis network point-to-point
```

IS-IS Fast Convergence

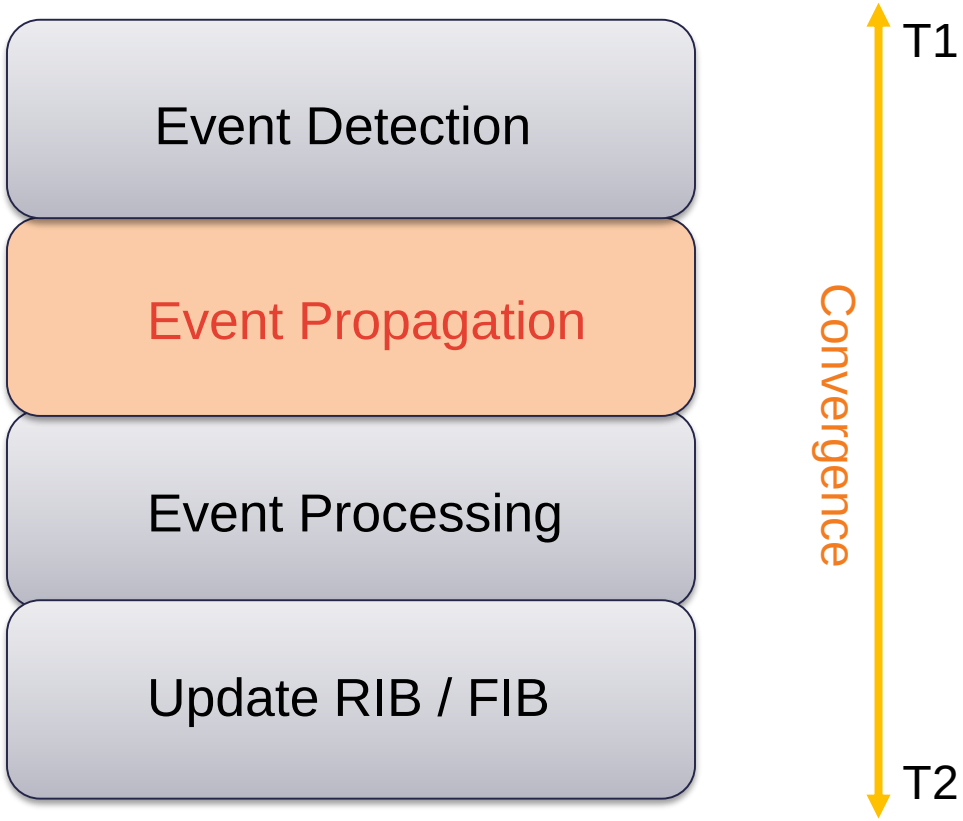

Event Propagation

LSP Fast Flooding

- The LSP needs to be flooded as fast as possible by the neighbor, when there is a change.
- When using short SPF-interval values for initial-delay, it may happen that SPF starts before the LSP who triggered SPF is flooded to Neighbors
- Specifically, flood the LSP who triggered SPF before starting SPF execution
- The fast-flood command ensures the first 5 LSPs that invoked SPF are flooded before running the SPF on the local router.

```
router isis
 fast-flood 15
```

LSP Flooding

LSP interval

- ISO 10589 states LSP flooding on a LAN should be limited to 30 LSP's per sec
- Default time between consecutive LSP's is a minimum of 33 milliseconds
- LSP pacing can be reduced in order to speed up end to end flooding
- Reduce the gap through: lsp-interval interface configuration command (msecs):

```
interface GigabitEthernet 1/0/0
 isis lsp-interval 10
```

Reduce the frequency and amount of flooding

- Reduce the amount of control traffic, conserving CPU usage for generation and refreshing of LSP's.
  - Do this by increasing the LSP lifetime to its limits.

```
router isis
 max-lsp-lifetime 65535
```

- Reduce the frequency of periodic LSP flooding of the topology, which reduces link utilization

```
router isis
 lsp-refresh-interval 65000
```

  - This is safe with the help of other mechanisms to guard against persistence of corrupted LSP's in the LSDB.

Ignore LSP errors

- Tell the IS to ignore LSP's with an incorrect data-link checksum, rather than purge them

```
router isis
ignore-lsp-errors
```

- Purging LSP's with a bad checksum causes the initiating IS to regenerate that LSP, which could overload the IS if continued in a cycle, so rather than purge them, ignore them.

IS-IS Fast Convergence

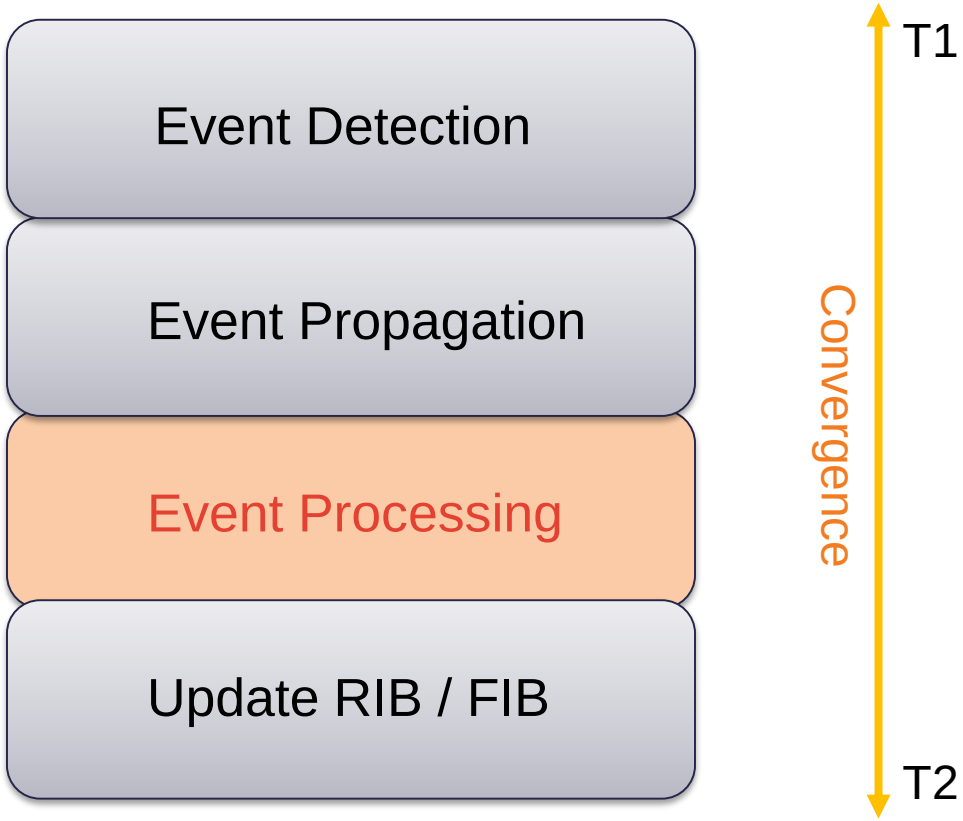

Event Processing

Throttling events

- IS-IS throttles the following events
  - SPF computation
  - PRC computation
  - LSP generation

- Throttling slows down convergence
- Not throttling can cause melt-downs
- Find a compromise...
- The scope is to react fast to the first events but, under constant churn, slow down to avoid collapse

Exponential Back-off Timer

- This mechanism dynamically controls the time between the receipt of a trigger and the processing of the related action.
  - In stable periods (rare triggers), the actions are processed promptly.
  - As the stability decreases (trigger frequency increases), the mechanism delays the processing of the related actions.

Exponential Back-off Timer

Throttling events

- These timers fine tunes three different events, which are a system of trigger and action

Trigger: Local LSP change ⟶ Action: Originate the new LSP and flood it

Trigger: LSP Database change and Tree Change ⟶ Action: Run SPF/iSPF

Trigger: LSP Database change but no tree change ⟶ Action: Run PRC

Exponential Back-off Timer

- The mechanism uses three parameters for all three events :
  - M (maximum) [s]
  - I (initial wait) [ms]
  - E (Exponential Increment) [ms]

```
router isis
  spf-interval M I E
  prc-interval M I E
  lsp-gen-interval M I E
```

Exponential Back-off Timer

Initial Wait, Maximum Time and Exponential Increment

- **Initial Wait (I) :**

Keeping Link failures in mind, set I = 1ms to trigger an SPF as soon as we receive a new LSP

With I = 1ms, convergence will be 5500ms faster in most cases, without any drawback (thanks to the dynamic adaptation provided by the exponential back-off algorithm)

Caveat : In some node failures (not all) and SRLG failures, we need several LSP's to be able to compute the right loop-free alternate path. If such cases are important, 'I' should be increased to several ten's of msec to ensure reception and flooding of these LSP's

Exponential Back-off Timer

Initial Wait, Maximum Time and Exponential Increment

- **Exponential timer (E) :**

Depends on how conservative : from 20msec to an average SPF time

If the first action took place and then a second trigger is received, the related action is scheduled to occur E after the previous action has been completed (timestamps are calculated at the end of each action). E is the Exponential Increment.

If the second trigger occurs in between the first trigger and the first action, obviously the first action is acted based on both triggers.

- **Maximum Time (M) :**

Again depends on how conservative - Default value looks fine except if frequency of bad/good news is high

Event Processing

SPF and RIB decoupled -PRC

- Run SPF (Dijkstra) only :
  - If any topology change (node, link)
  - Recompute SPT and the RIB
- Run PRC (Partial Route Calculation):
  - If only an IP prefix changes
  - keep the SPT
  - just update the RIB for the nodes whose prefixes have changed

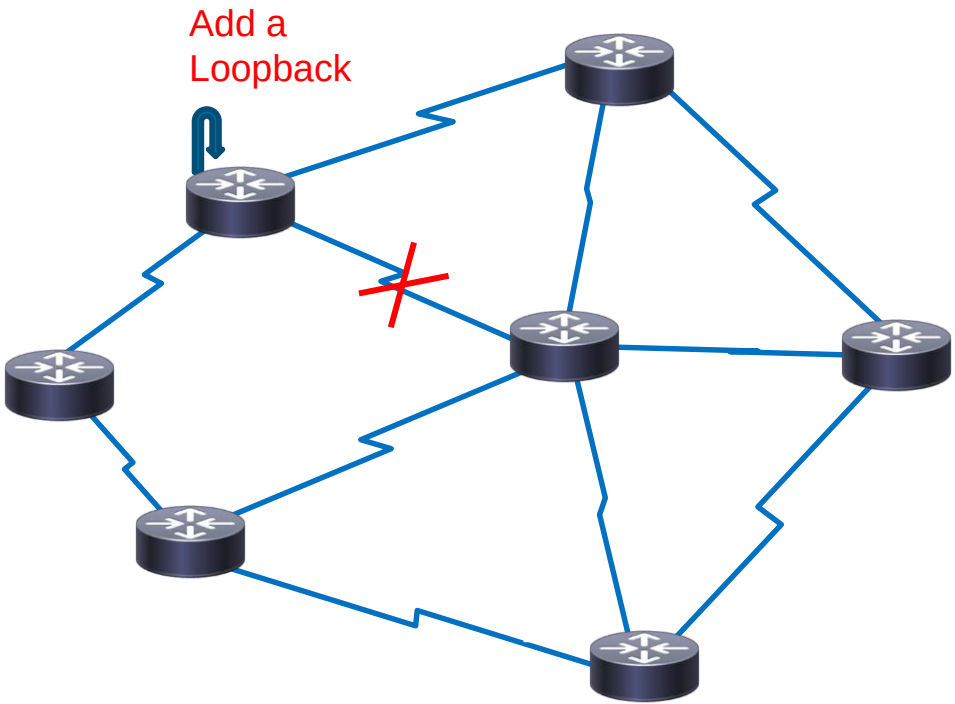

Enable iSPF (incremental SPF)

- iSPF in the long run, reduces CPU demand because SPF calculations are run only on the affected changes in the SPT

```
router isis
  ispf [level-1 | level-2 | level-1-2]
```

- On L1-L2 routers, enable iSPF at both levels. Configure the timer (seconds) for ispf to start, after the command has been entered into the configuration

```
ispf level-1-2 60
```

IS-IS Fast Convergence

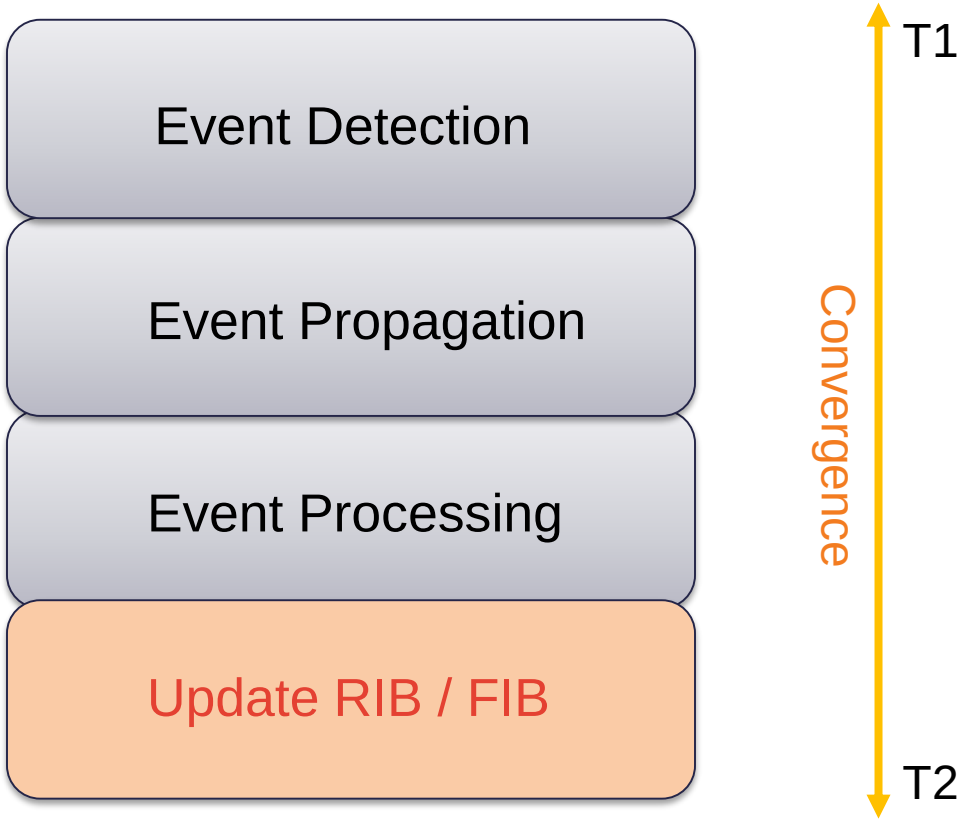

Update the RIB / FIB tables

LSDB, RIB and FIB

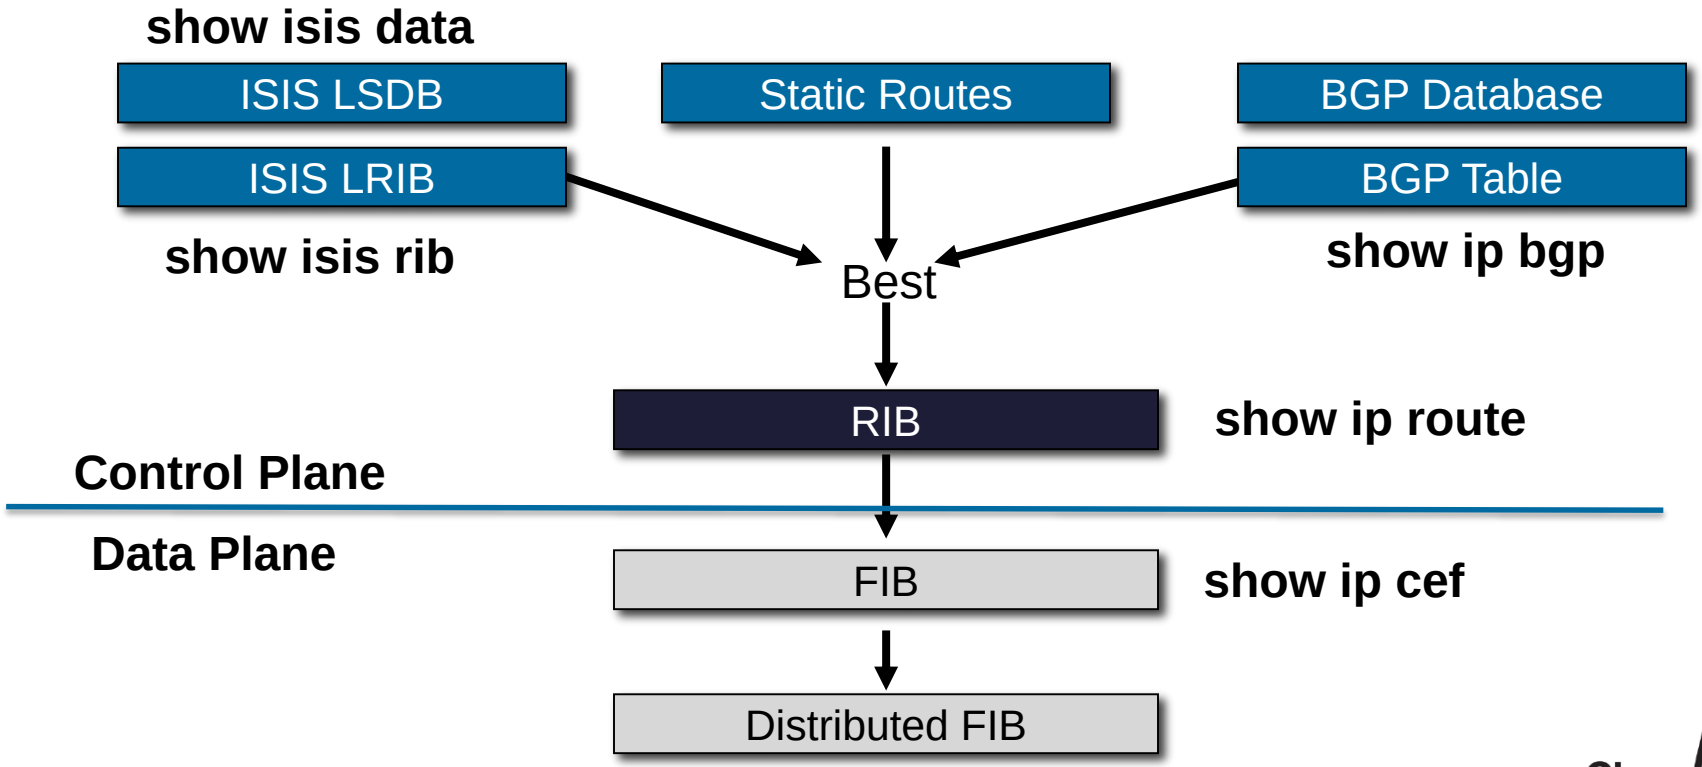

Optimize RIB/FIB update

- RIB update:
  - linear function of the number of prefixes to update
  - worst-case = function of the total number of prefixes to update

- Design principle: optimize the number of prefixes in the IGP
  - At the extreme, a designer could recognize that the only important prefixes that should be present in the IGP are those tracking explicit content destinations (a subnet with servers) and BGP next-hops.
  - All the other prefixes only track links interconnecting routers and this information may be advertised in I-BGP.

RIB: Limit the number of prefixes

- Limit the number of ISIS prefixes to the minimum to scale. There are two options :
  - Exclude the connected interfaces manually – better control, works for small scale

```
int GigabitEthernet4/1
  ip router isis
  no isis advertise-prefix
```

  - Just advertise loopback's prefix , which is passive, works for large scale

```
router isis
  advertise-passive-only
```

RIB/FIB update

Prioritize IS-IS Local RIB

- Customization via the ISIS tag
  - Important prefixes (VoIP subnet, Video sources) may be tagged and hence will converge first
- Worst-case RIB update duration for important prefixes is now bounded by
  - the number of important prefixes * RIB update time per prefix
  - rather than the total number of prefixes * RIB update time per prefix.
- As in most cases, the number of important prefixes is significantly smaller than the total number of prefixes, this functionality is extremely useful and is a significant fast-convergence gain.

Prefix Prioritization

- Prefix Prioritization is a key differentiator
  - CRITICAL: IPTV SSM sources
  - HIGH: Most Important PE's
  - MEDIUM: All other PE's
  - LOW: All other prefixes
- Prefix prioritization customization is generally required for CRITICAL and HIGH

RIB: Local RIB and prefix prioritization

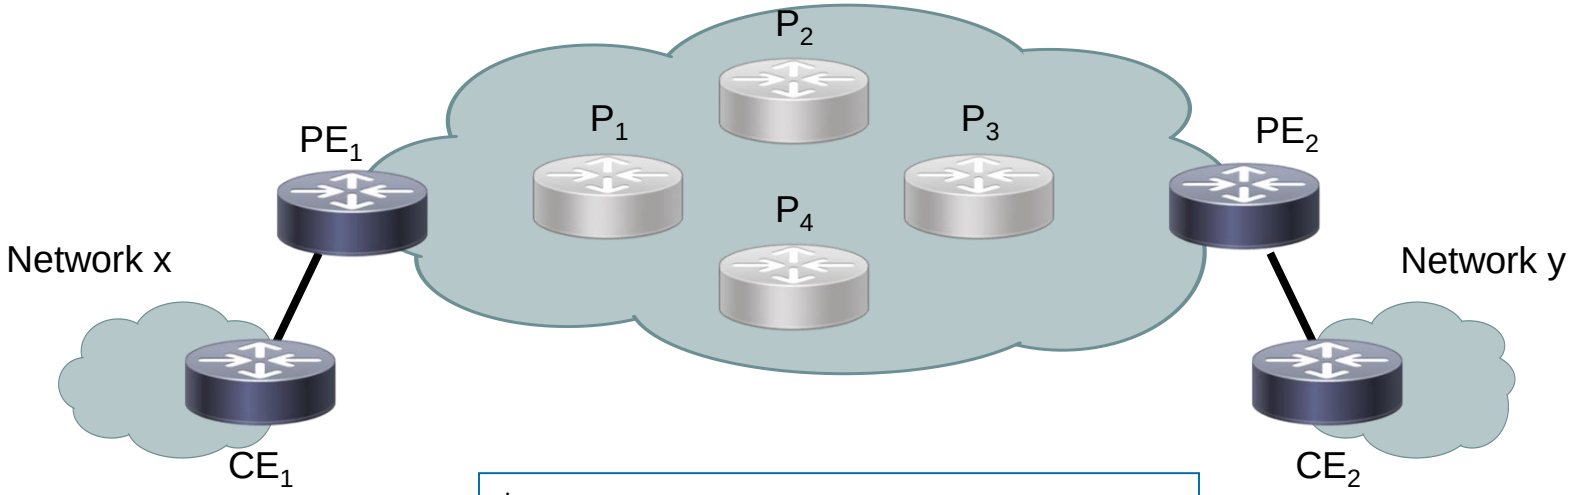


```
!
interface loopback0
 ip router isis
 isis tag 17
!
router isis
  ip route priority high tag 17
```

Agenda

- ISIS Overview
  - CLNS, L1/L2 Routing, Best Practices
- ISIS for IPv6
  - Single Topology, Multi-Topology
- ISIS in the Backbone
  - Area Design, Fast Convergence Features
- ISIS at the Edge
  - BGP and MPLS Considerations
- ISIS at the Access / Aggregation
  - Route Leaking, Traffic Engineering and IP FRR

Interaction with BGP

set-overload-bit

- Mechanism used by IS-IS Networks in order to decrease the data loss associated with deterministic black-holing of packets during transient network conditions
- “set-overload-bit” condition can be used by a router in a transient condition to tell other routers not to use itself as a transit node
- Typically when IS-IS is up but BGP may not have had time to fully converge or even MPLS not up yet
- Better stabilization in the network

```
router isis
 set-overload-bit [on-startup[<timeout>|wait-for-bgp]]
```

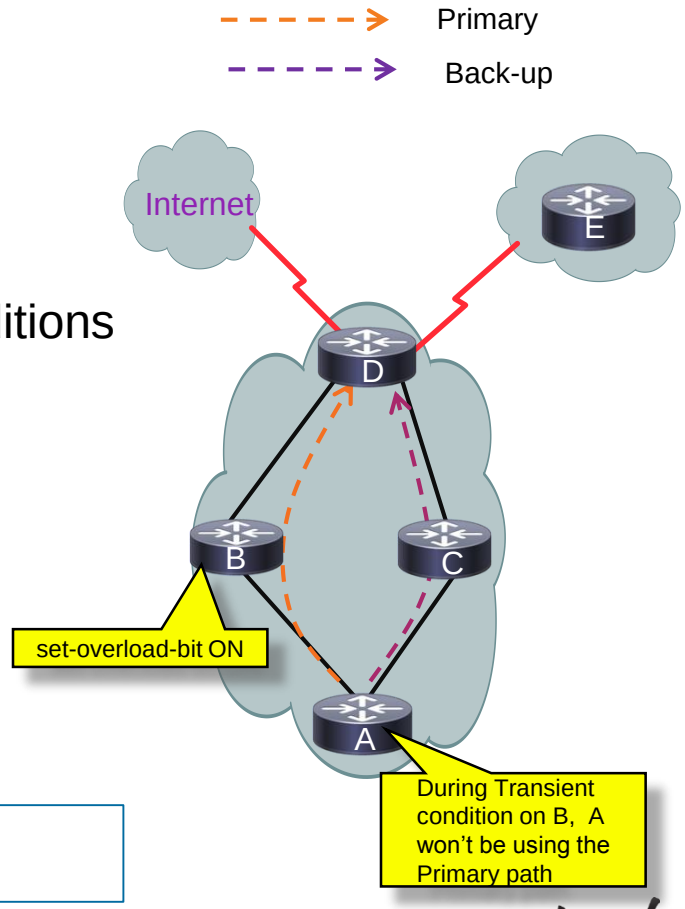

IS-IS and LDP

- Problem statement
  - If IGP selects the link before the LDP labels are available any MPLS-VPN (L2/L3) traffic is lost until the labels are ready

- Solution
  - LDP session protection
  - LDP/IGP synchronization

LDP session Protection

- A loopback-to-loopback session is automatically established upon LDP neighbor detection on a local interface
- This session survives link failure for <seconds> (default: indefinitely) and hence ensures that the labels of the neighbors are still present when the link comes back up
- This requires redundant path between the two nodes, which can be non-direct (typically the case in SP backbone)

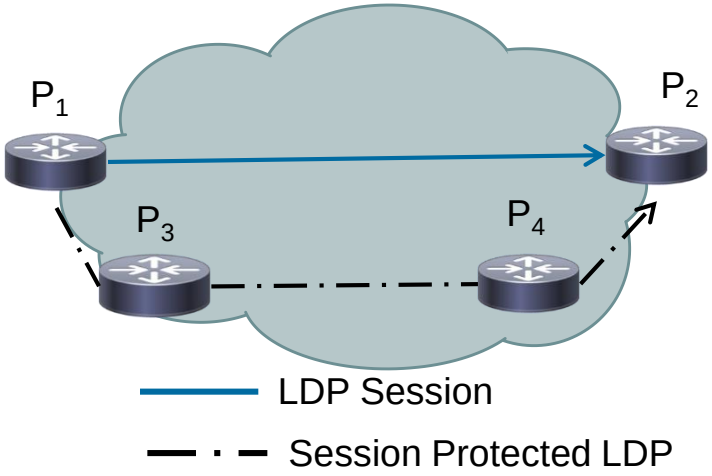


```
mpls ldp session protection [ for <acl> | duration <seconds> | vrf <name> ]
```

LDP/IGP Sync

- LDP sessions and traffic loss:
  - When an adjacency goes UP, traffic might start flowing across the link, even before the LDP session is UP.
  - If an LDP session goes DOWN, forwarding might continue over the broken link, instead of using a better path.

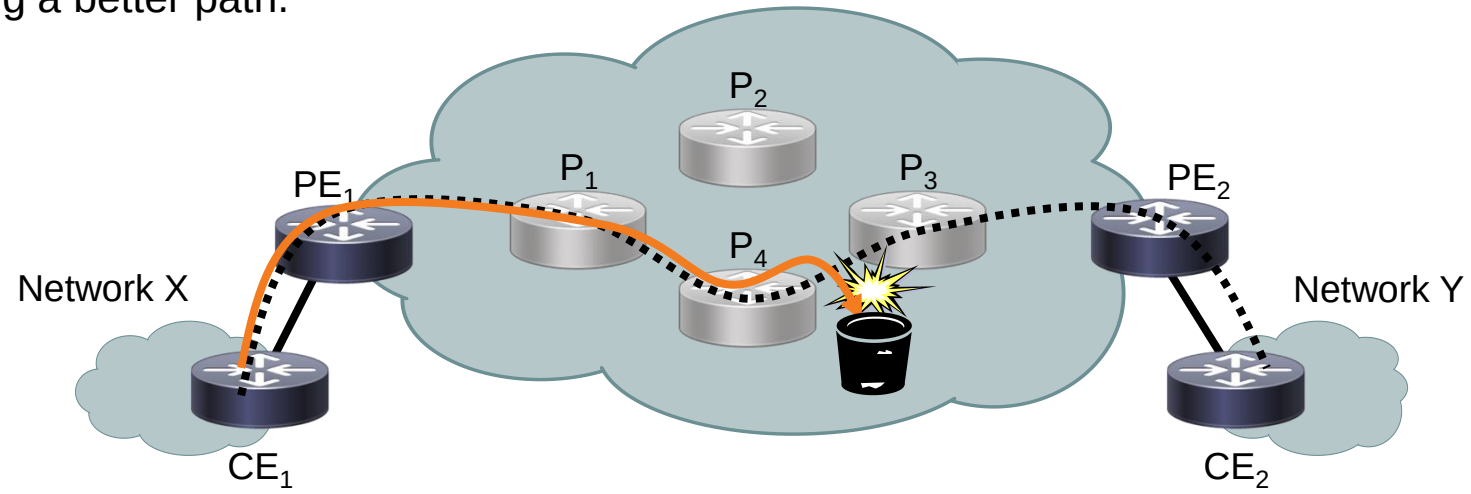

LDP/IGP Sync

- Keep the IGP State Synchronized with LDP session State

```
router isis
 mpls ldp sync
```

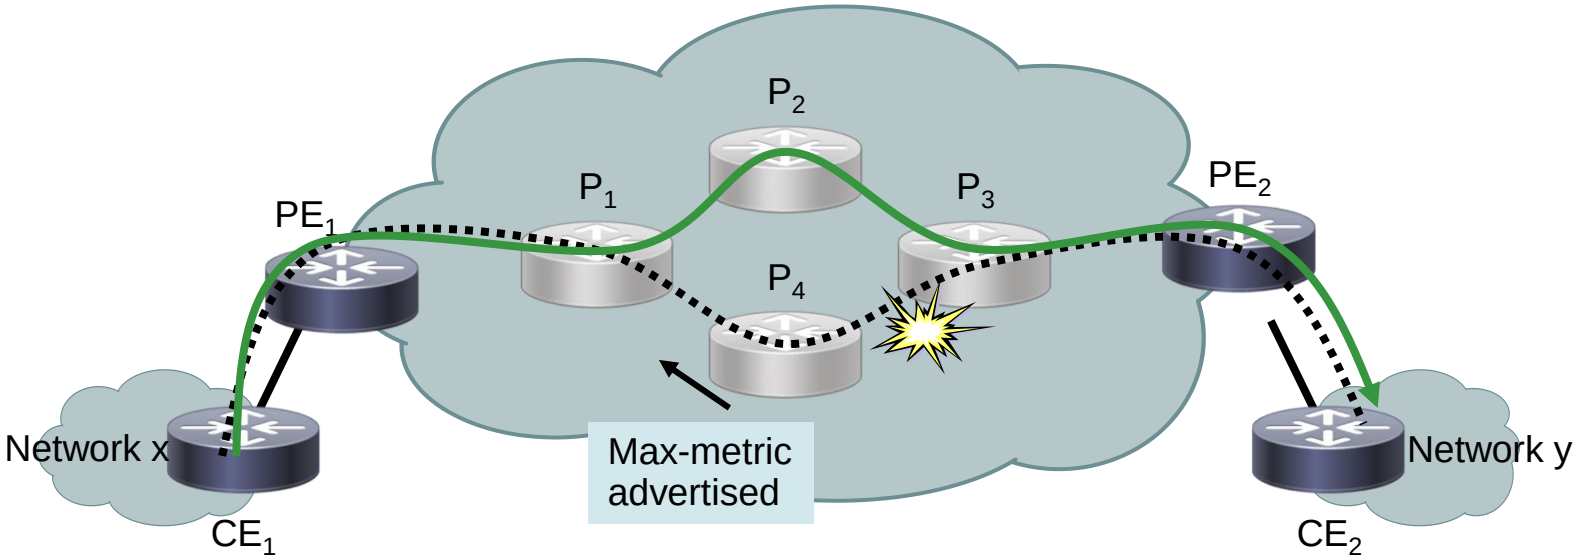

NSR and NSF (Graceful Restart)

- Intra-chassis recovery mechanisms with dual supervisors
- The IS-IS NSF feature offers two modes:
  - IETF = NSF (Non-stop forwarding)

```
router isis
  nsf ietf
```

  - Cisco = NSR (Non-stop Routing)

```
router isis
  nsf cisco
```

- Software and platform support is limited, so check whether your particular platform/code supports this. Also, deploy only if it's necessary

NSR and NSF (Graceful Restart)

- IETF mode (NSF):
  - With IETF, Operation between peer devices based on a proposed standard. But neighbors need to be NSF-aware
  - After the switchover, neighbor routers provide adjacency and link-state information to help rebuild the routing information following a switchover.

- Cisco mode (NSR) :
  - Neighbors do not need to be NSF aware
  - Using the Cisco configuration option, full adjacency and LSP information is saved, or “check-pointed”, to the redundant supervisor engine.
  - Following a switchover, the newly active supervisor engine maintains its adjacencies using the check-pointed data, and can quickly rebuild its routing tables.

NSF and Hello Timers

- When NSF/SSO is included in the design, a good objective is to avoid losing the hello adjacency during a valid switch-over.
- This ensures an IS going through a control plane switchover continues to forward traffic as if nothing happened
- In most scenarios, testing has indicated that the “hold down” should not be configured to less than 4 seconds to achieve this.
- In networks with only P2P links or BFD, IGP will re-converge as soon as the interface goes down or a failure happens, NSF will not work.

Agenda

- ISIS Overview
  - CLNS, L1/L2 Routing, Best Practices
- ISIS for IPv6
  - Single Topology, Multi-Topology
- ISIS in the Backbone
  - Area Design, Fast Convergence Features
- ISIS at the Edge
  - BGP and MPLS Considerations
- ISIS at the Access / Aggregation
  - Route Leaking, Traffic Engineering and IP FRR

L1-L2 Router at Edge of POP

Route-Leaking

- It is recommended to configure the L1-L2 routers at the edge of the pop with route-leaking capabilities
- Leak BGP next-hops and summarize physical link
- Hence the L1 routers will be able to take the right exit/entry router based on the metric of the leaked IP-prefix
  - Optimal Inter-Area Routing
- Ensure ‘metric-style wide’ is configured when leaking routes
  e.g. MPLS-VPN (PEs Loopback Reachability and LSP binding)

ISIS LFA Fast Reroute

LFA – Loop Free Alternate

- Backup Path is pre-computed using LFA mechanism so router can very rapidly switch when a failure is detected without further computation
- Traffic is re-routed while IGP converges
- Backup Paths are computed AFTER the primary path and so do not delay normal convergence
- A fast detection mechanism is required to trigger the forwarding engine to switch from the primary path to the backup path (BFD…)

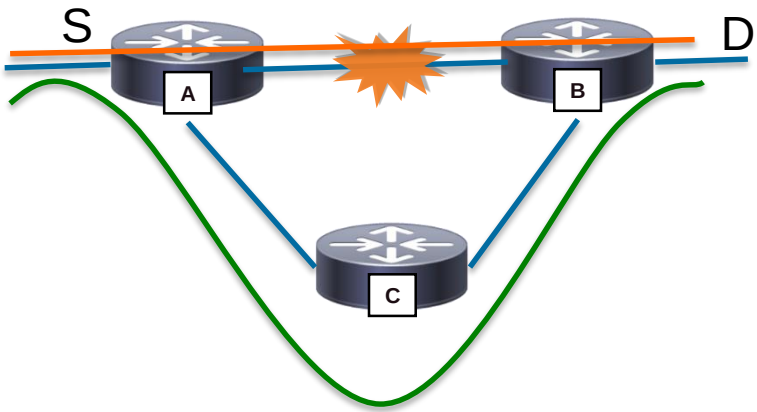

LFA Conditions

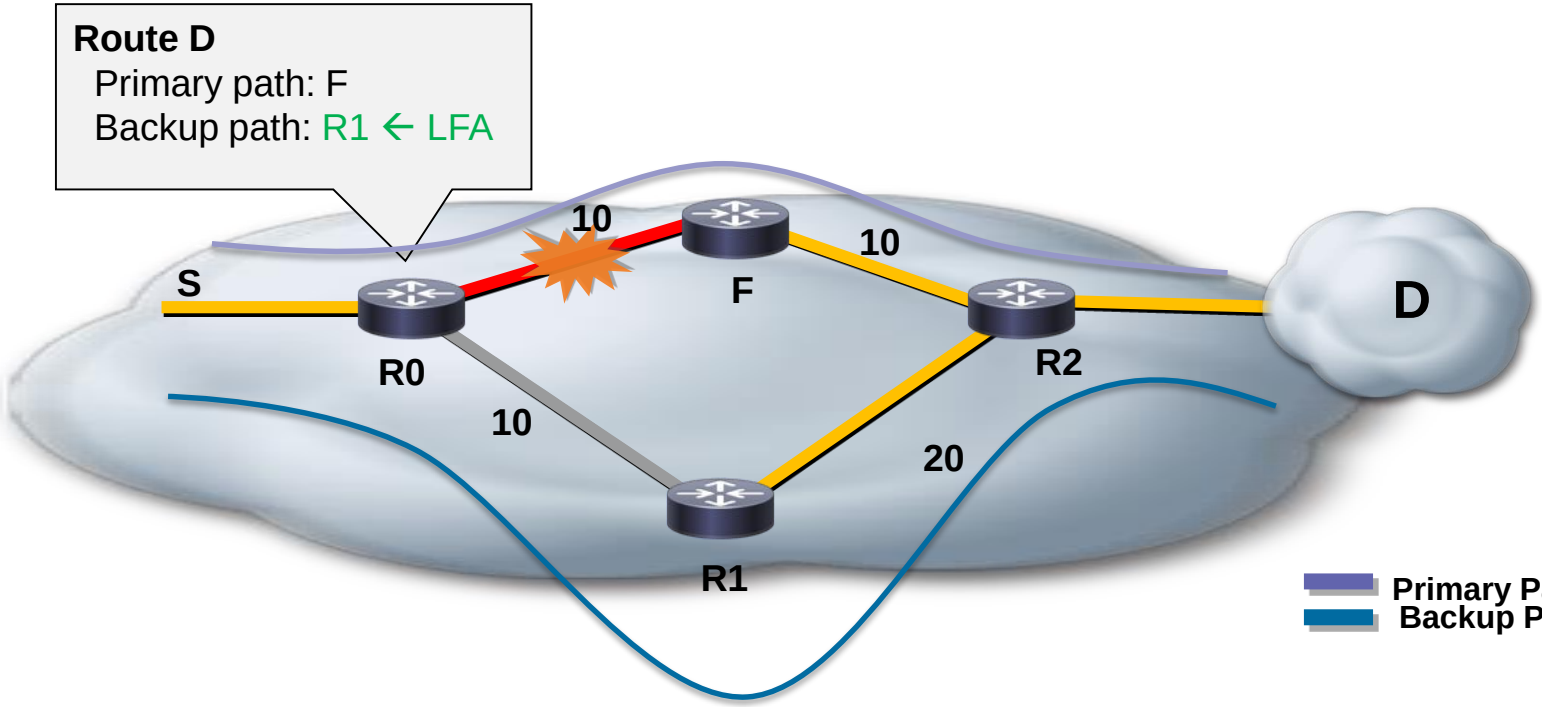

Conditions with no LFA

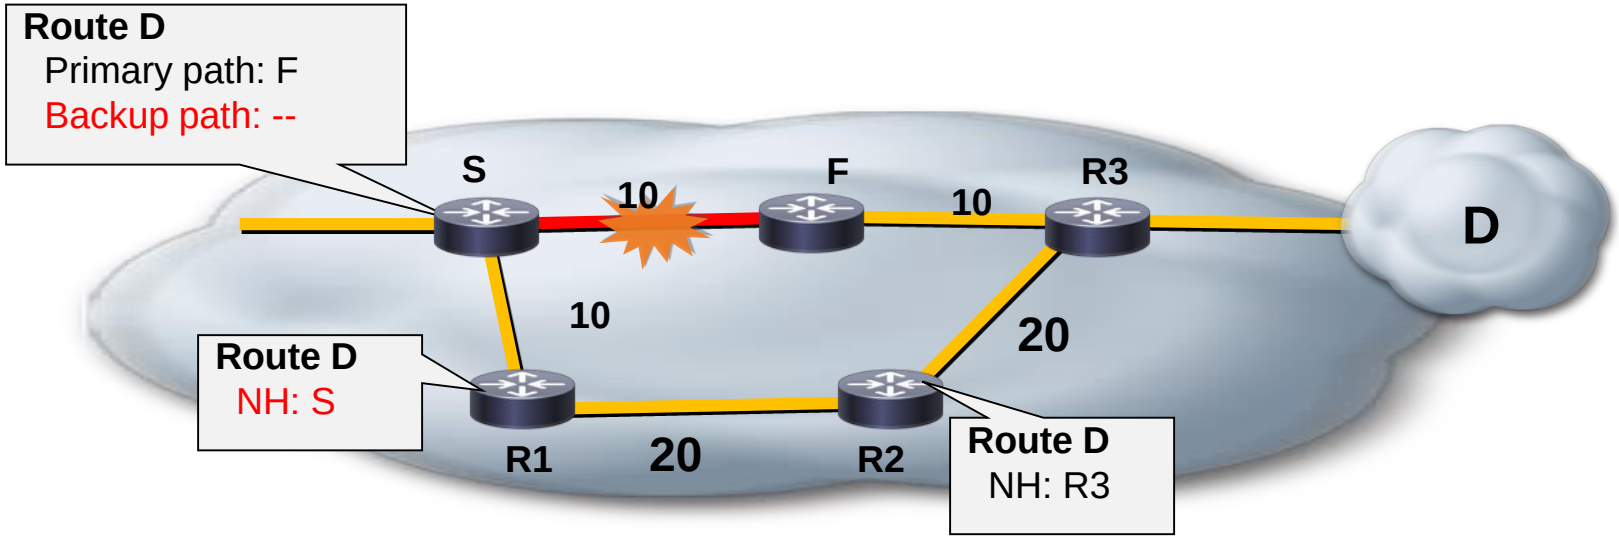

ISIS – Enabling LFA on IOS

- By default, LFA computation is disabled
- To enable LFA computation

```
router isis
  fast-reroute per-prefix {level-1 | level-2} {all | route-map <route-map-name>}
```

- Default action if enabled :
  - LFA computations is enabled for all routes
  - FRR is enabled on all supported interfaces

ISIS – FRR Using Route Maps

Protecting BGP next-hops using interface tags

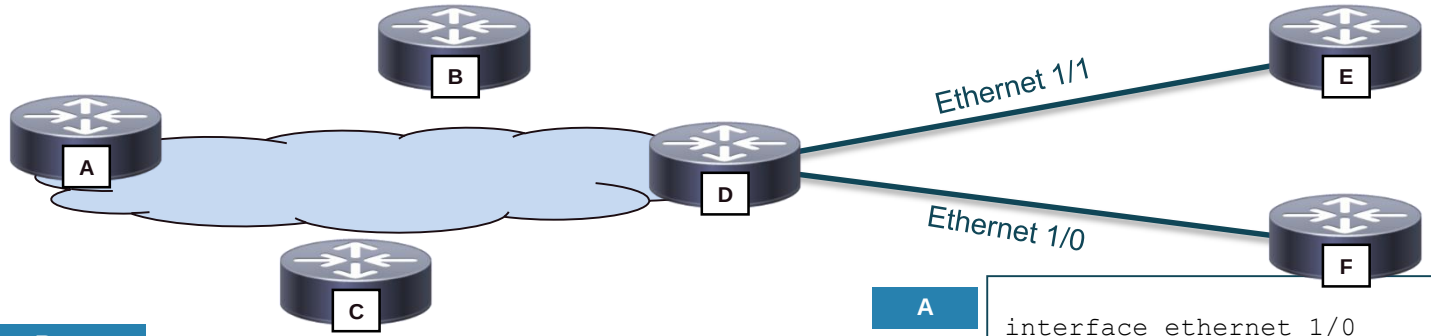

Other Routers

```
router isis
 net 47.0004.004d.0001.0001.c11.1111.00
 fast-reroute per-prefix level-2 route-map ipfrr-include
!
route-map ipfrr-include
  match tag 17
```

Route tags are 4 bytes long and flooded with LSAs in sub-TLV 1 of TLV 135

A

```
interface ethernet 1/0
 ip address 10.1.1.1 255.255.255.0
 ip router isis
 isis tag 17
interface ethernet 1/1
 ip address 172.16.1.1 255.255.255.0
 ip router isis
 isis tag 17
router isis
 net 49.0001.0001.0001.0001.00
 fast-reroute per-prefix level-2
```

Summary: What Have We Learned?

- Deploying IS-IS from Scale, Convergence and Ease of troubleshooting standpoint
- Considerations with single Area / Multi-Area design
- Deploying IPv6 with IS-IS and migration techniques
- Techniques to achieve fast convergence in different parts of the network
- IS-IS features to optimize operations with BGP and MPLS
- Best practices and recommendations for every segment

Reference Configuration with Best Practices on IOS and IOS-XR

Cisco live!

IS-IS Configuration on IOS

```
router isis
 net 10.0000.0000.0010.00
 is-type level-2-only
 advertise passive-only
 metric-style wide
 fast-flood
 ip route priority high tag 10
 set-overload-bit on-startup wait-for-bgp
 max-lsp-lifetime 65535
 lsp-refresh-interval 65000
 spf-interval 2 50 100
 hello padding
 nsf cisco | ietf
 fast-reroute per-prefix level-2 all
 redistribute isis ip level-2 into level-
1 distribute-list 199
 passive-interface Loopback
 bfd all-interfaces
!
```

```
interface TenGigabitEthernet3/2
 ip address 192.168.1.1 255.255.255.252
 ip router isis
 bfd interval 200 min_rx 200 multiplier 3
 isis circuit-type level-2-only
 isis network point-to-point
 no isis advertise prefix
 isis tag 10
 isis mesh-group
!
```

IS-IS Configuration on IOS-XR

```
router isis DEFAULT
 set-overload-bit on-startup wait-for-bgp
 is-type level-2-only
 net 10.0000.0000.0009.00
 nsf cisco | ietf
 lsp-refresh-interval 65000
 max-lsp-lifetime 65535
 address-family ipv4 unicast
  metric-style wide
  fast-reroute per-prefix priority-limit critical
  fast-reroute per-prefix priority-limit critical
  spf-interval maximum-wait 2000 initial-wait 50
 secondary-wait 150
  advertise passive-only
 !
 interface Loopback0
  passive
 !
```

```
interface TenGigE0/0/0/0
  bfd fast-detect ipv4[8]
  mesh-group 1
  point-to-point
  hello-padding sometimes
  address-family ipv4 unicast
  !
!
interface TenGigE0/0/0/2
  point-to-point
  address-family ipv4 unicast
  !
 !
!
```

Configuring IS-IS for MPLS TE on IOS-XR

```
mpls traffic-eng tunnels
!
interface TenGigabitEthernet0/1/0
 ip address 172.16.0.0 255.255.255.254
 ip router isis
 mpls traffic-eng tunnels
 mpls traffic-eng attribute-flags 0xF
 mpls traffic-eng administrative-weight 20
 ip rsvp bandwidth 100000
!
router isis
 net 49.0001.1720.1625.5001.00
 is-type level-2-only
 metric-style wide
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng level-2
 passive-interface Loopback0
!
```

Configuring IS-IS for MPLS TE on IOS-XR

```
router isis DEFAULT
 is-type level-2-only
 net 49.0001.1720.1625.5129.00
 address-family ipv4 unicast
  metric-style wide
  mpls traffic-eng level 2
  mpls traffic-eng router-id Loopback0
 !
 interface Loopback0
  passive
  address-family ipv4 unicast
  !
 !
 interface TenGigE0/0/0/0
  address-family ipv4 unicast
  !
 !
!
```

```
rsvp
 interface TenGigE0/0/0/0
  bandwidth 100000
 !
!
mpls traffic-eng
 interface TenGigE0/0/0/0
  admin-weight 5
  attribute-flags 0x8
 !
!
```

Recommended Sessions

- BRKRST-3020 - Advanced - IP LFA (Loop-Free-Alternative): Architecture and Troubleshooting
- BRKRST-2124 - Introduction to Segment Routing
- BRKRST-3122 - Segment Routing: Technology and Use-cases
- BRKRST-3123 - Segment Routing for IPv6 Networks
- BRKRST-2022 - IPv6 Routing Protocols Update
- BRKRST-2336 (EIGRP), 2337 (OSPF) – Deployment in Modern Networks
- BRKRST-3371 – Advances in BGP
- BRKMPL-3101 - Advanced Topics and Future Directions in MPLS
- LTRSPG-2500 - L2VPN over IOS-XE and IOS-XR: Configuration, Deployment and Troubleshooting
- BRKRST-2044 - Enterprise Multi-Homed Internet Edge Architectures

Call to Action

- Visit the World of Solutions for
  - Cisco Campus – Enterprise Networks, Service Provider
  - Walk in Labs – MPLS / Routing labs
  - Technical Solution Clinics - Routing
- Lunch time Table Topics
- DevNet zone related labs and sessions
- Recommended Reading: for reading material and further resources for this session, please visit www.pearson-books.com/CLMilan2015

Complete Your **Online Session Evaluation**

- Please complete your online session evaluations after each session. Complete 4 session evaluations & the Overall Conference Evaluation (available from Thursday) to receive your Cisco Live T-shirt.

- All surveys can be completed via the Cisco Live Mobile App or the Communication Stations

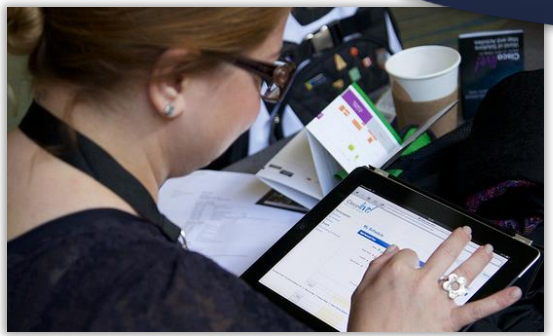

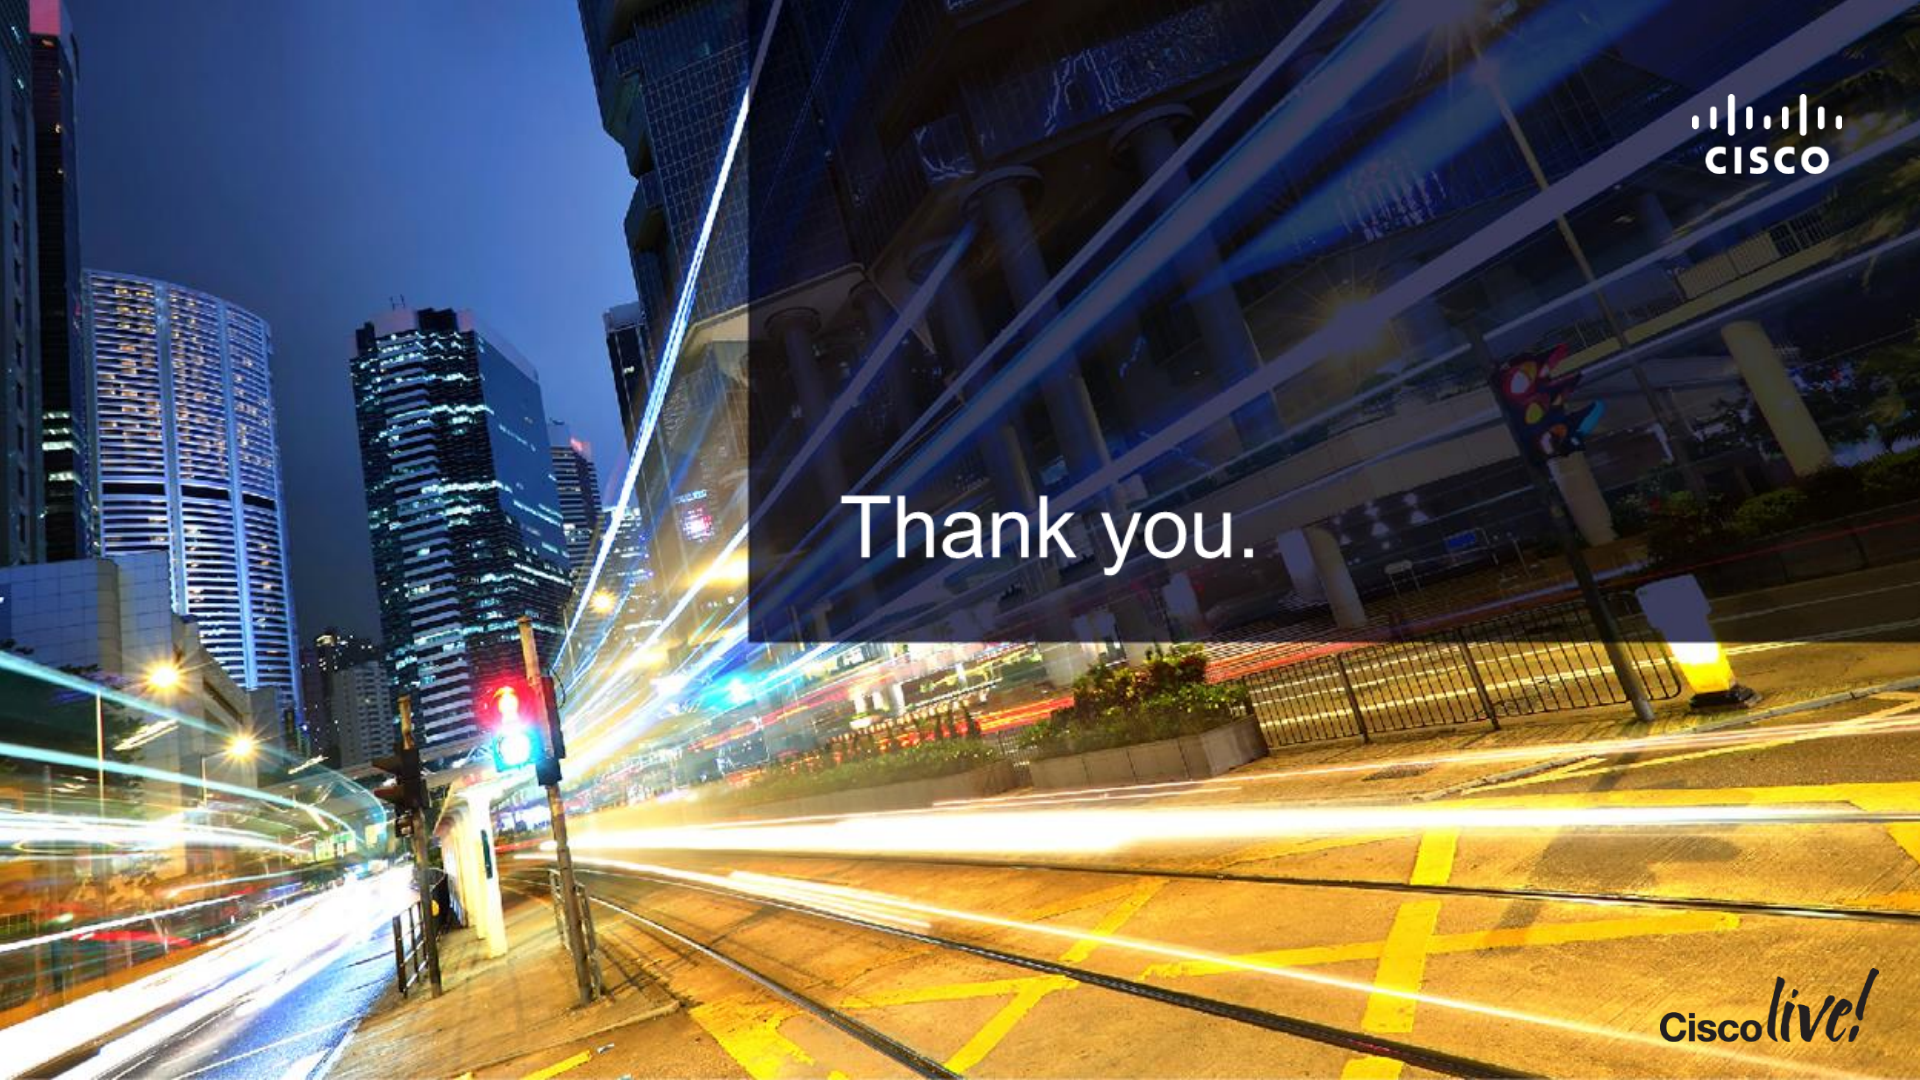
Thank you.
CISCO
Cisco live!


CISCO