



# 2 NETWORKERS



A man in a white shirt and red tie is holding a large red pipe over a colorful landscape. The pipe is arched over the landscape, and the man is standing on a blue hill. The background is a mix of yellow and blue.

# The Integrated ISIS routing protocol

**Stefano Previdi**  
sprevidi@cisco.com



# Agenda

- **Link-state protocol fundamentals**
- **Overview of IS-IS**
- **Areas and levels**
- **NSAPs and LSP identifiers**
- **CLNS routing principles**
- **LSP Flooding**

# Agenda

- **IP routing specifics**
- **Design issues**
- **New features**
- **Future enhancements**



# Link-state protocol fundamentals

# About link-state protocols

- In a link-state protocol, the network can be viewed as a jigsaw puzzle
- Each jigsaw piece holds one router
- Each router creates a packet which represents its own jigsaw piece

**This packet is called a Link State PDU (LSP)**

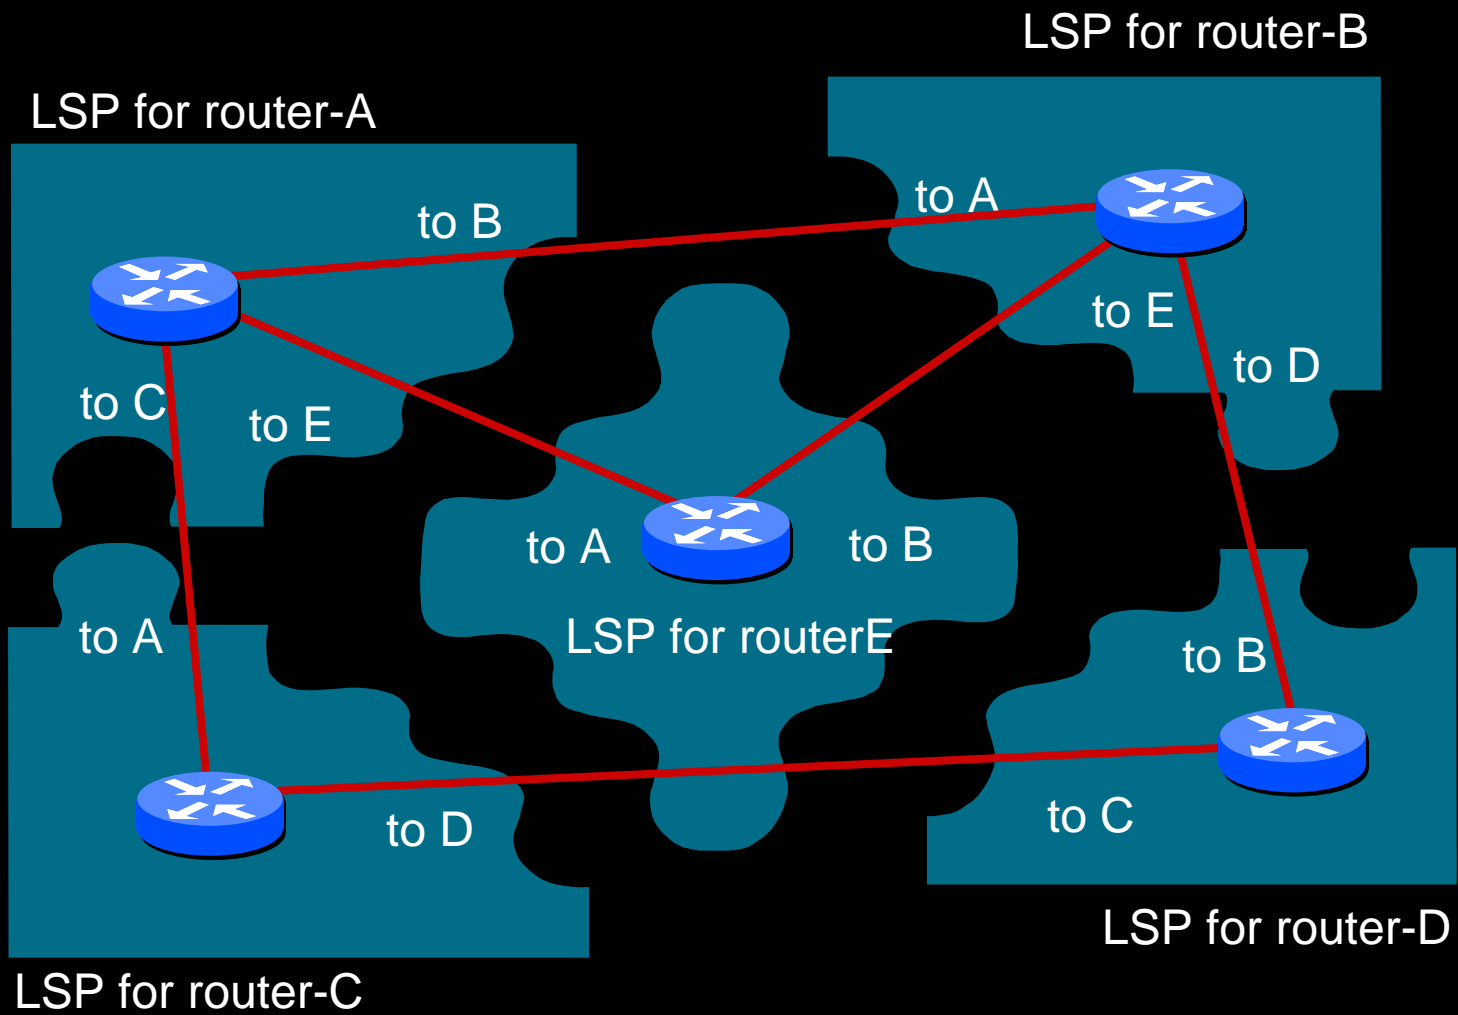
# About link-state protocols

- **These packets are flooded everywhere**
- **Therefore each router receives all pieces of the jigsaw puzzle**
- **Each routers compute SPF algorithm to put the pieces together**

**Input: all jigsaw puzzle pieces (LSPs)**

**Output: Area or network topology tree  
Shortest Path Tree**

# The jigsaw puzzle



# All routers have same view

- **All routers exchange all LSPs**  
via a reliable flooding mechanism
- **All routers store all LSPs in a so-called link-state database (LSPDB)**  
separate from the routing table (RIB)  
all routers should have exactly the same LSPDB, but different RIBs

# What to do with LSPs ?

- Each router 'composes the jigsaw puzzle' by executing Dijkstra's Shortest Path First algorithm (SPF)


the topology is calculated as a Shortest Path Tree (SPT), with itself as root

each router computes a different SPT

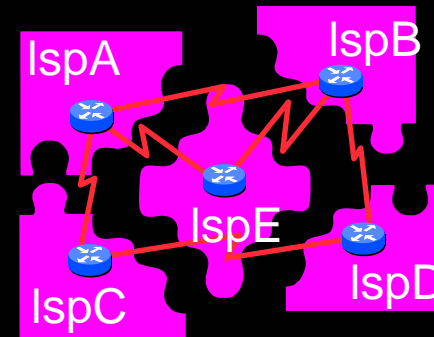
- From the SPT the RIBs are calculated

# All routers have same LSPDB


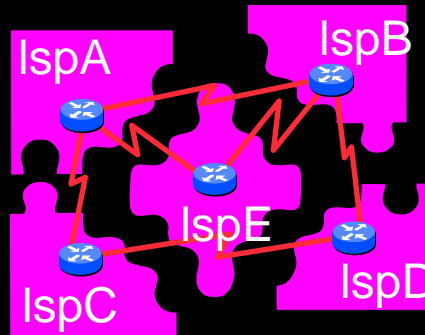
## RouterA's LSPDB



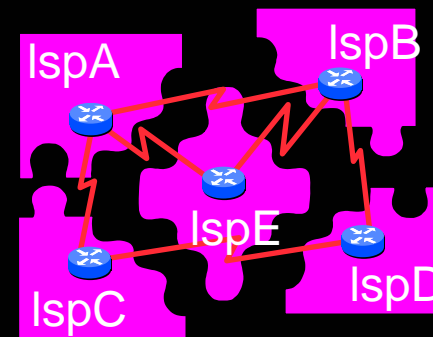
## RouterB's LSPDB



## RouterE's LSPDB



## RouterC's LSPDB



## RouterD's LSPDB



# Brief overview of IS-IS

# What is IS-IS ?

- **IS stands for Intermediate System**
- **IS is “OSI speak” for router**
- **IS-IS is the Intermediate System to Intermediate System intra-domain routing protocol**
- **IS-IS was defined in 1992 in the ISO/IEC recommendation 10589**

# IS-IS for IP routing

- **IS-IS was designed for OSI routing**
- **IS-IS is easily extendable**
- **Extensions for IP routing in rfc1195**
- **Also called “Integrated IS-IS” or “Dual IS-IS”**
- **Easy to extend for other protocols**  
**mainly IPv6**

# Only 4 types of IS-IS packets

- **IS-IS Hello packet (IIH)**
- **Link State Packet (LSP)**
- **Partial Sequence Number Packet (PSNP)**
- **Complete Sequence Number Packet (CSNP)**
- **Packets are sometimes called Protocol Data Units (PDU in OSI)**

# IS-IS Hello PDUs

- **Also called IIHs**
- **Used for maintaining adjacencies**
- **Different on p2p links and LANs**
- **Different from ISHs and ESHs (ESIS)**
- **IIHs are padded to full MTU size**

# Link State PDUs

- **Called LSPs**
- **Contains all info about one router**  
adjacencies, connected IP prefixes, OSI  
endsystems, area addresses, etc.
- **One LSP per router (plus fragments)**
- **One LSP per LAN network**

# Sequence Number PDUs

- **Partial (PSNP) and Complete (CSNP)**
- **Used when flooding the LSPDB**
- **PSNPs are like ACKs on p2p links**
- **CSNPs are used for LSPDB synchronization over LANs**
- **CSNP are also used to sync LSPDB over new p2p adjacencies**

# Pseudonodes and network LSAs

- **For SPF, the whole network must look like a collection of nodes and point-to-point links**

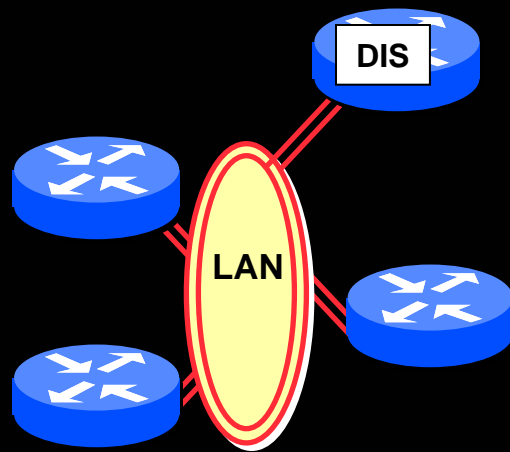
**Multi-access networks are different**

- **Assume a virtual node for the LAN**

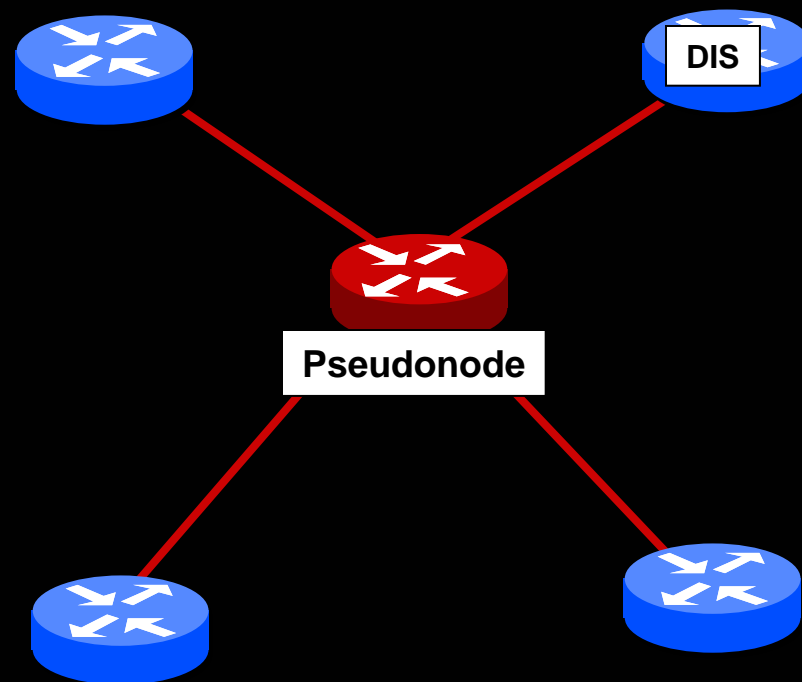
**this virtual node is called pseudonode. It is not a real router, but just an extra LSP in the LSPDB**

# A pseudonode on a LAN

Physical view



Logical view



# Who creates the pseudonode

- **Created by Designated Router (DIS)**
- **Not by Backup Designated Router. In IS-IS there even is no BDR**
- **The DIS reports all LAN neighbors in the pseudonode LSP**

**With metric 0**

- **All LAN routers report connectivity to the pseudonode in their LSPs**

# Identifying nodes

- **In IS-IS SystemIDs are 6 bytes. Nodes are identified by 7 bytes.**

**A normal node (non-pseudonode) is identified by 6 bytes systemID plus a zero**

**A pseudonode is identified by the systemID of the DIS, plus 1 byte from the circuitID of the interface of the DIS**



# Areas and levels

# Hierarchy

- **IS-IS has 2 layers of hierarchy**  
the backbone is called level-2  
areas are called level-1
- **Same algorithms apply for L1 and L2**
- **A router can take part in L1 and L2**  
inter-area routing (or inter-level routing)

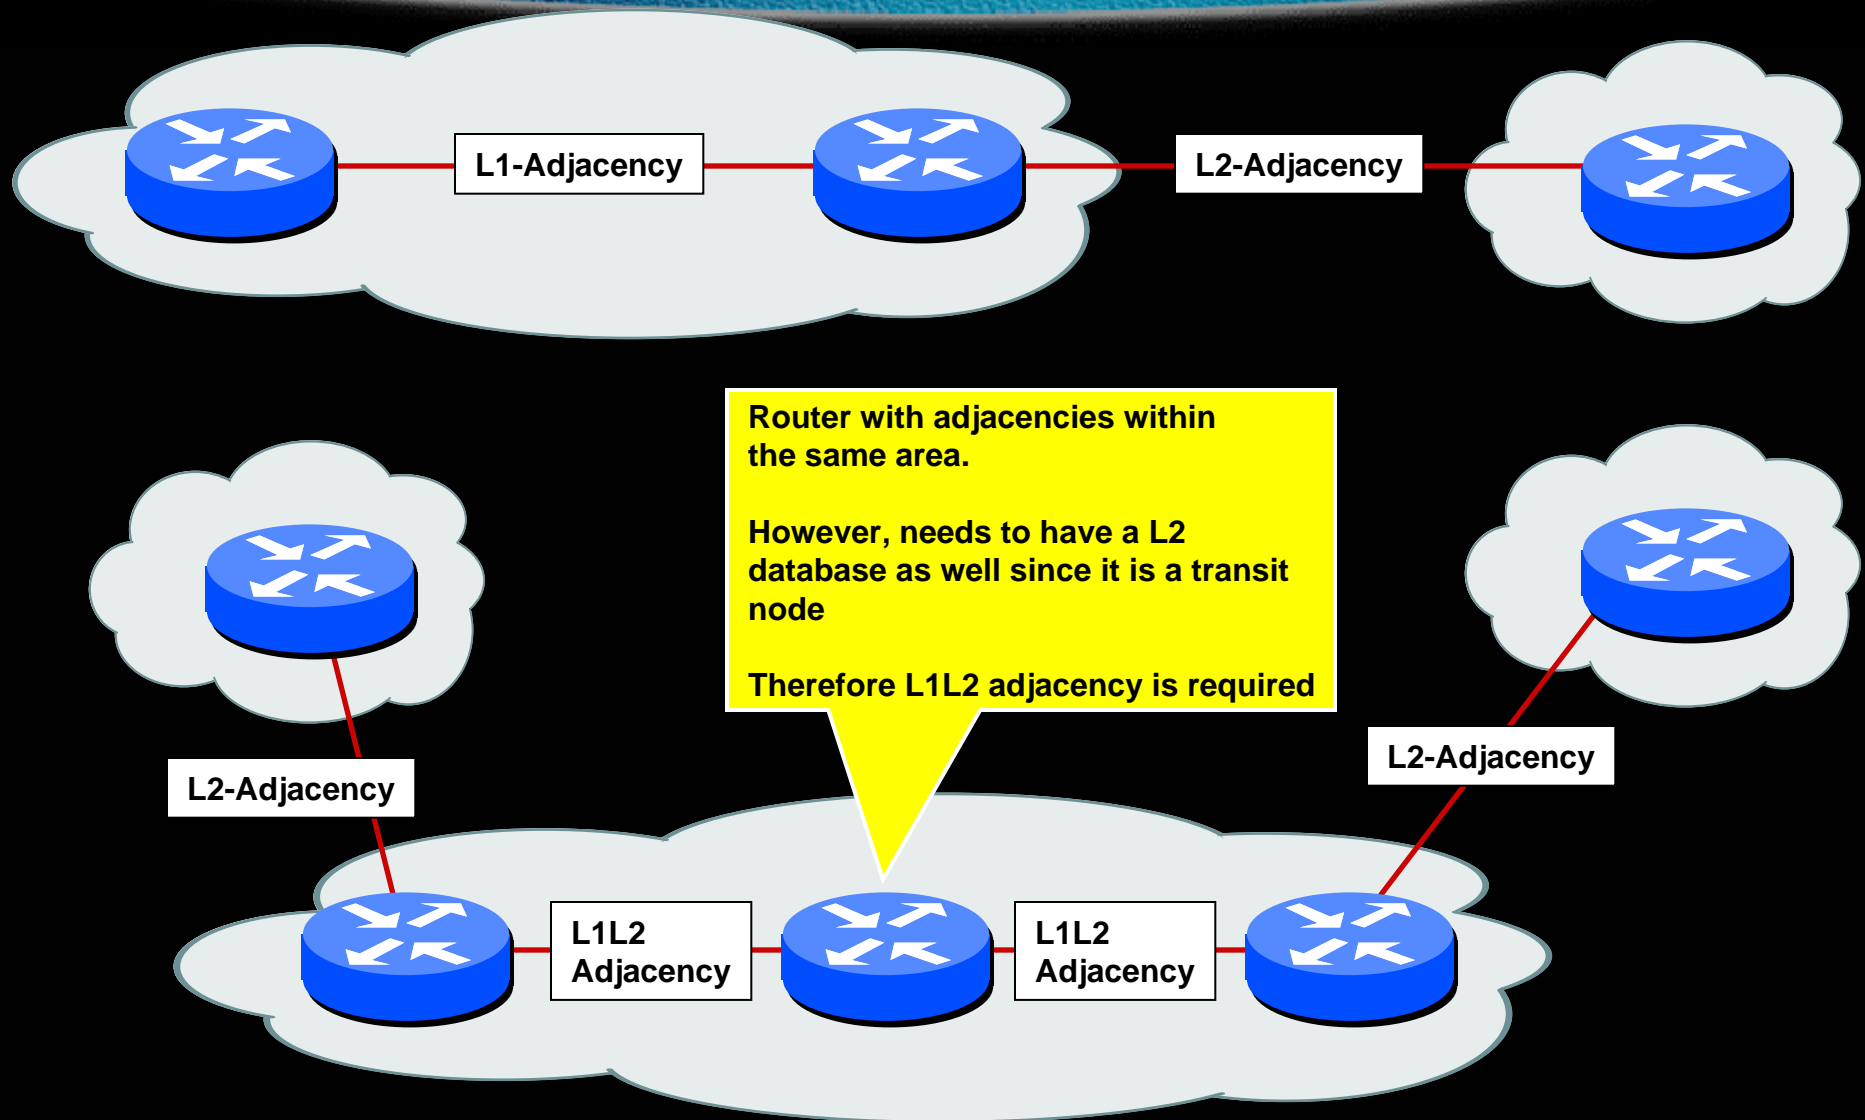
# Level-1 Routers

- **Neighbors only in the same area**
- **L1 has information about own area**
- **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**

# Level-2 routers

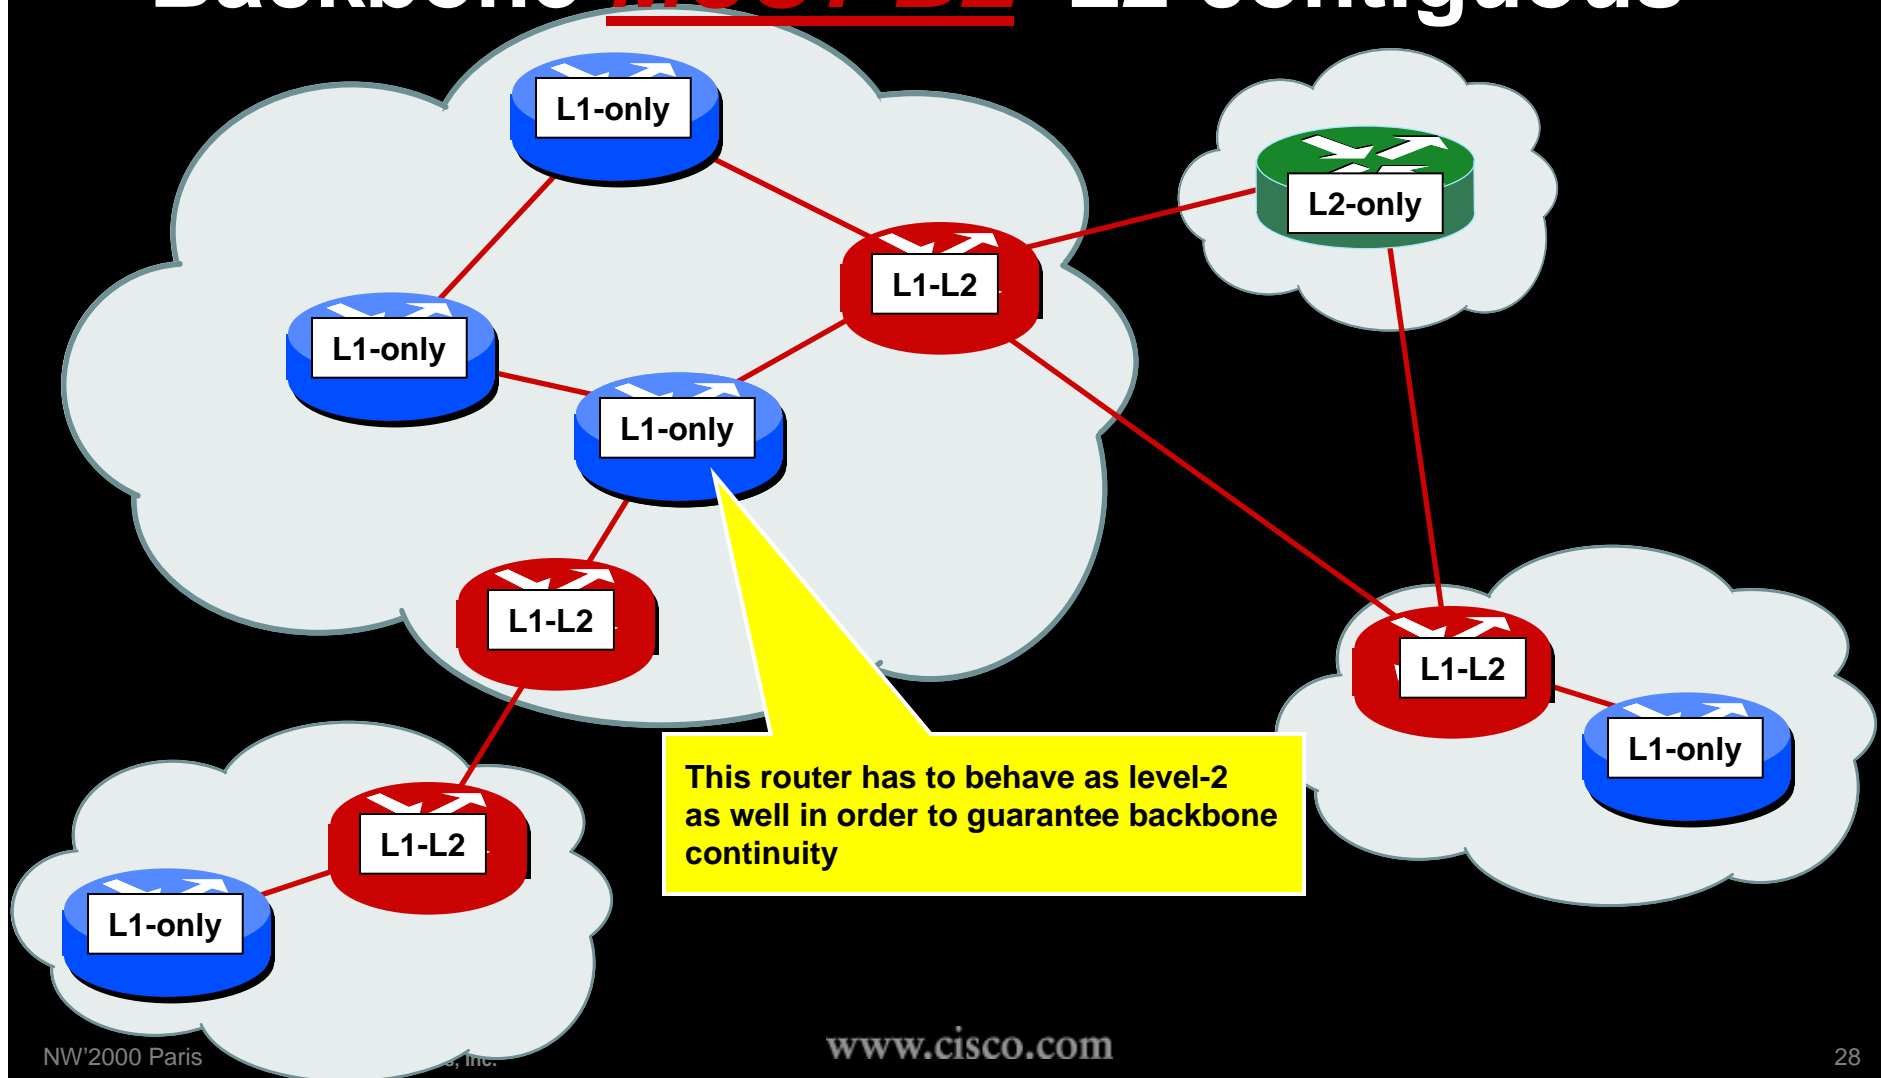
- **May have neighbors in other areas**
- **L2 has information about L2 topology**
- **L2 has information what L1 destinations are reachable and how to reach them via the L2 topology**
- **L2 routers often do also L1 routing  
so called L1L2 routers**

# Adjacency levels



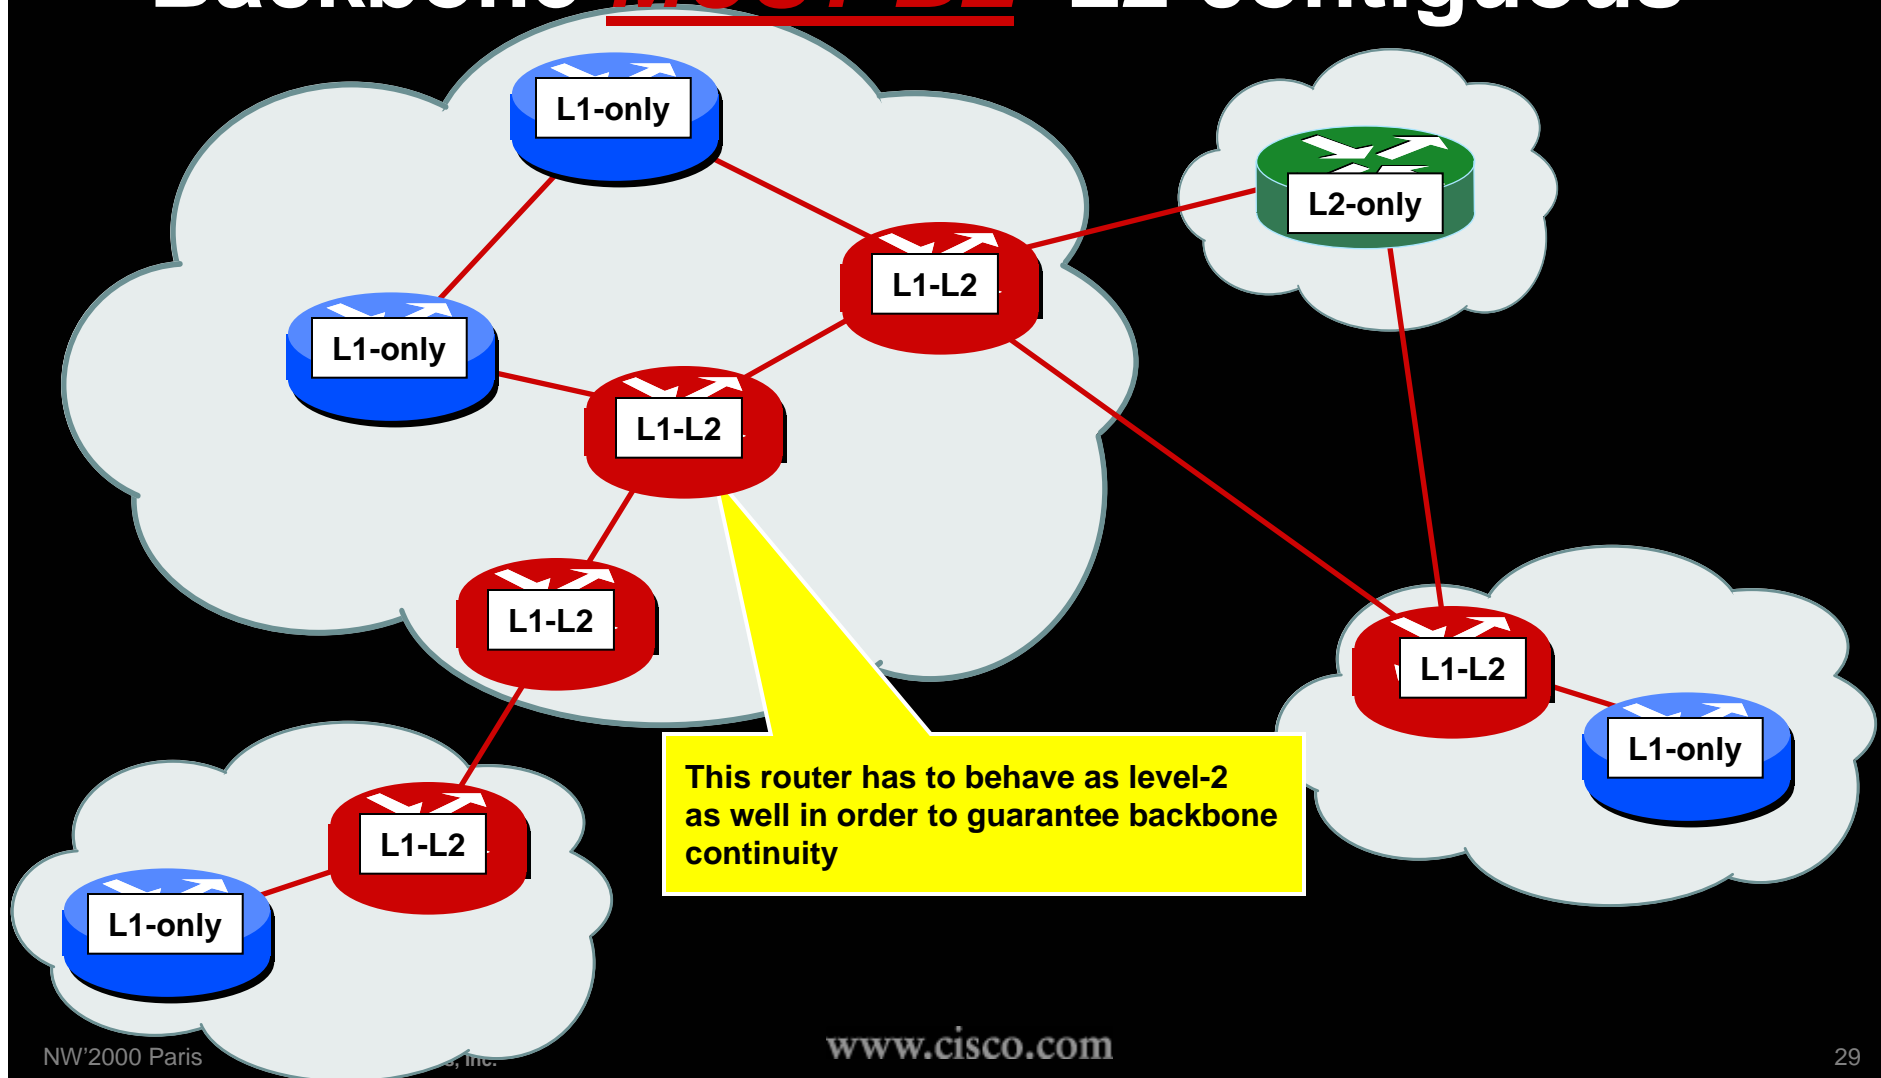
# Level-1, Level-2 & Level-1-2 Routers

- Backbone **MUST BE** L2 contiguous



# Level-1, Level-2 & Level-1-2 Routers

- Backbone **MUST BE** L2 contiguous






# NSAPs and LSPids

# NSAPs and Addressing

- **Network Service Access Point**
- **The NSAP is the network layer address for CLNS packets**
- **One NSAP per box, not per interface**
- **SNPA means SubNetwork Point of Attachment, which is the layer2 or MAC address**

# How do I read an NSAP ?



- **An NSAP consists of 3 parts  
area-address, systemID and n-selector**
- **Total length between 8 and 20 bytes**

**example: 49.0001.0000.0000.0007.00**

# NETs versus NSAPs

- A NET is an NSAP with n-selector 0
- A NET implies the routing layer of the IS itself (no transport layer)
- On routers we always deal with NETs

**We haven't implemented TP4 (or another transport layer)**

# Do I need an NSAP if I want to use IS-IS for IP routing ?

- **Yes, still needed for IP routing only**
- **Area address is like OSPF area nr**
- **SystemID is like an OSPF routerID**  
**LSP identifier is derived from systemID**

# Creating unique systemIDs

- **SystemID is 6 bytes**
- **Start numbering 1, 2, 3, 4 .... etc**
- **Convert your loopback IP address**

**192.31.231.16 -> 192.031.231.016 ->**  
**systemID                    1920.3123.1016**

# Creating area addresses

- If you do CLNS routing, request an official NSAP prefix
- If you do just IP routing, use AFI 49
- AFI 49 denotes private address space like network 10.0.0.0 in IP
- Just number your areas 49.0001... 49.0002....., 49.0003,... etc

# LSP Identifier

- **LSP identifier consists of 3 parts**
  - **Source ID**  
SystemID of router or DIS (if pseudonode)
  - **Pseudonode ID**  
Router LSP = zero, Pseudonode LSP = non-zero
  - **LSP number**  
Fragmentation number

• **Example: 00c0.0040.1234.02-00**

SystemID                      PN-ID                      Frag-Nr



# CLNS routing principles

# OSI protocol family

- **CLNS is datagram delivery protocol like bare IP service actually called CLNP**
- **ESIS is like ARP, ICMP, HSRP, IRDP between routers and hosts**
- **IS-IS and ISO-IGRP are the IGPs**

# OSI protocol family

- **IS-IS is not encapsulated in CLNS and not encapsulated in IP (yet)**
- **Encapsulated directly in layer2**
- **Protocol family is OSI**  
**usually values like 0xFE or 0xFEFE**  
**(ppp uses 0x0023 and 0x8023)**

# Level-1 routing

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

# Area addresses

- **An area address is like a summary**
- **All L1 ISs and ESs in an area must have NSAPs that start with the same area-address**
- **L1L2 routers advertise their area-addresses to L2 routers in other areas**
- **Multiple area-addresses possible**

# Level-2 Routing

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

# Level-2 Routing

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

# The Backbone

- **A router can't tell whether it is a transit node**

**Therefore the cisco default is to be L1L2**

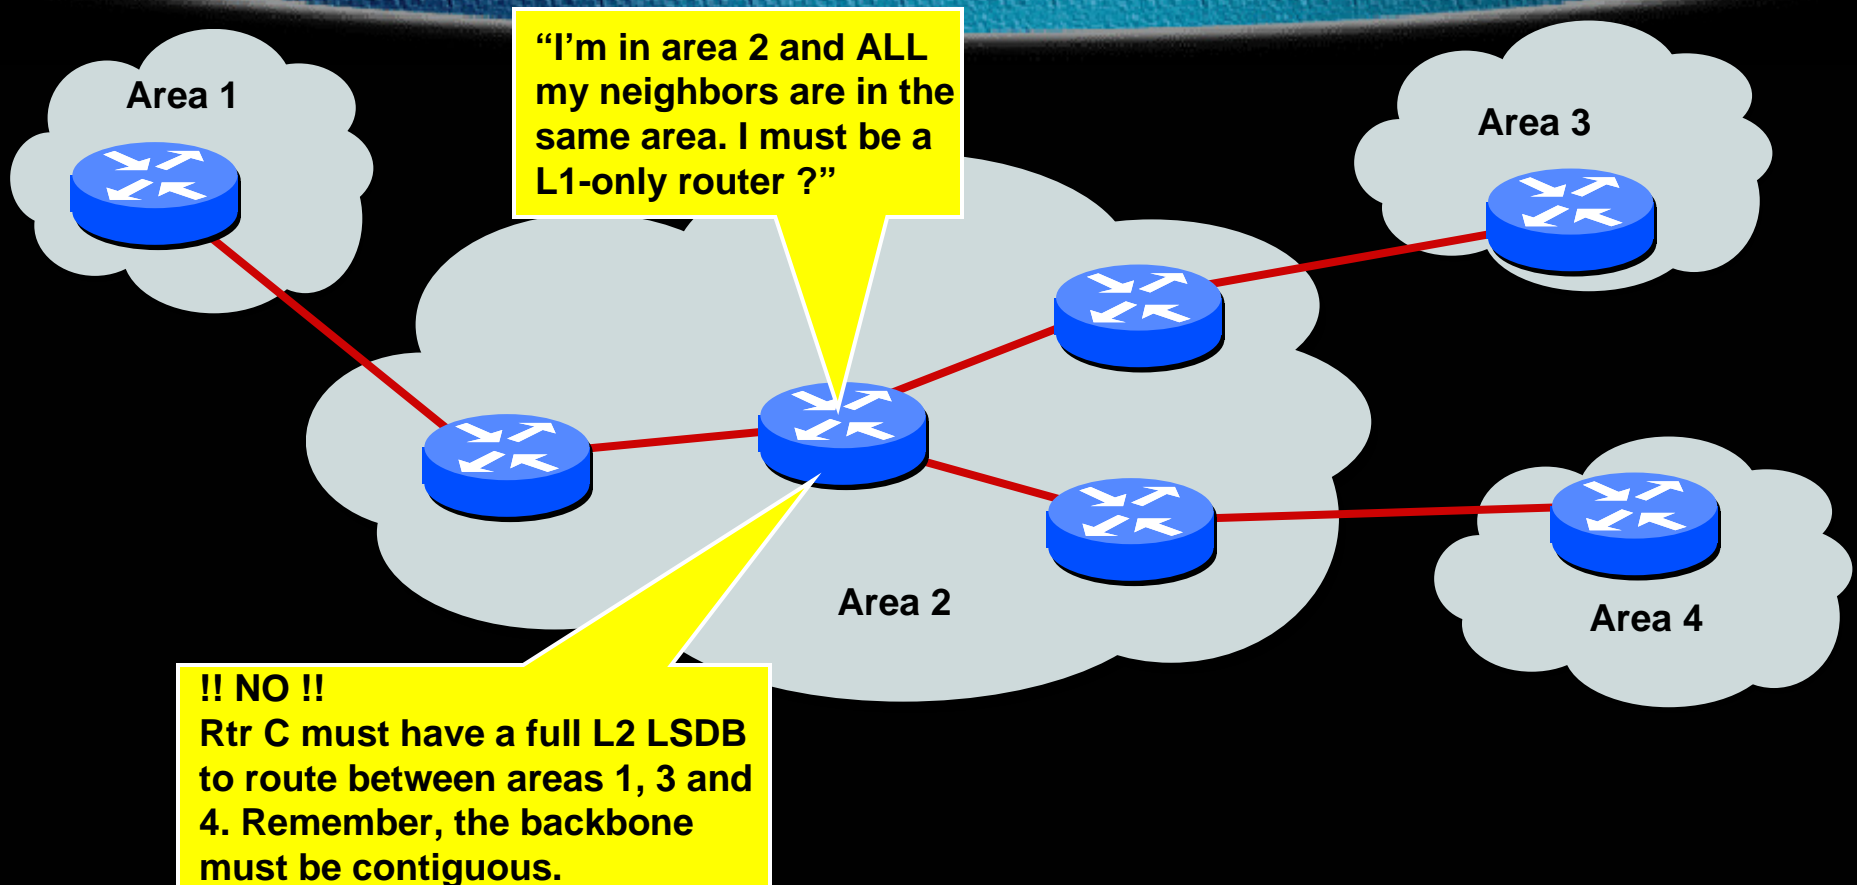
**This will make the backbone larger than necessary**

**So always configure L1-only or L2-only when possible**

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


**Especially with ISIS for IP**

# Can an IS determine its level ?



**ISIS router cannot determine if they need to be L1 or L1L2  
Therefore By default all cisco routers will behave as L1L2**

# 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 towards the L1L2 with ATT-bit set

- Shortest metric to the L1L2 who sets the ATT bit wins

A man in a white shirt and dark tie is holding a long, curved pipe in a blue-tinted environment. The pipe is arched over him, and he is looking up at it. The background is a textured, blue surface.

# LSP flooding

# Why do we need flooding

- **All routers generate an LSP**
- **All LSPs need to be flooded to all routers in the network**
  - if LSPDB is not synchronised, routing loops or blackholes might occur
- **IS-IS' two components are the SPF computation and reliable flooding**

# What triggers a new LSP ?

- **When something changes ...**

**Adjacency came up or went down**

**Interface up/down (connected IP prefix !)**

**Redistributed IP routes change**

**Inter-area IP routes change**

**An interface is assigned a new metric**

**Most other configuration changes**

**Periodic refresh**

# What to do with a new LSP ?

- **Create new LSP, install in your own LSPDB and mark it for flooding**
- **Send the new LSP to all neighbors**
- **Neighbors flood the LSP further**

# Basic flooding rules

- **When receiving an LSP, compare with old version of LSP in LSPDB**
- **If newer:**
  - install it in the LSPDB**
  - Acknowledge the LSP with a PSNP**
  - Flood to all other neighbors**
  - Check if need to run SPF**

# Basic flooding rules

- **If same age:**
  - Acknowledge the LSP with a PSNP**
- **If older:**
  - Acknowledge the LSP with a PSNP**
  - Send our version of the same LSP**
  - Wait for PSNP**

# Sequence number

- **Each LSP (and LSP fragment) has its own sequence number**
- **When router boots, set seqnr to one**
- **When there is a change, the seqnr is incremented, a new version of the LSP is generated with the new seqnr**
- **Higher seqnr means newer LSP**

# Remaining lifetime

- **Used to age out old LSPs**
- **Periodic refresh needed to keep stable LSPs valid**
- **IS-IS counts down from 1200 sec to 0**  
**we allows to start at 65535 sec (18.7h)**
- **When lifetime expires, the LSP is purged from the network**  
**Header with lifetime = 0 is flooded**


# Flooding on a P2P Link

→  
LSP  
id=x seqnr=22

Receives LSP  
id=x seqNr=22  
It's new. Put it in  
the LSPDB

Now flood it:  
Set SRM bit.  
Send over p2p.

Received ack  
Clear SRMbit



→  
LSP  
id=x seqnr=22

Received it. Local  
copy has seqNr = 21.  
So the received one is  
newer. Install it in LSDB.  
Acknowledge it. Maybe  
flood further.

←  
PSNP  
id=x seqnr=22

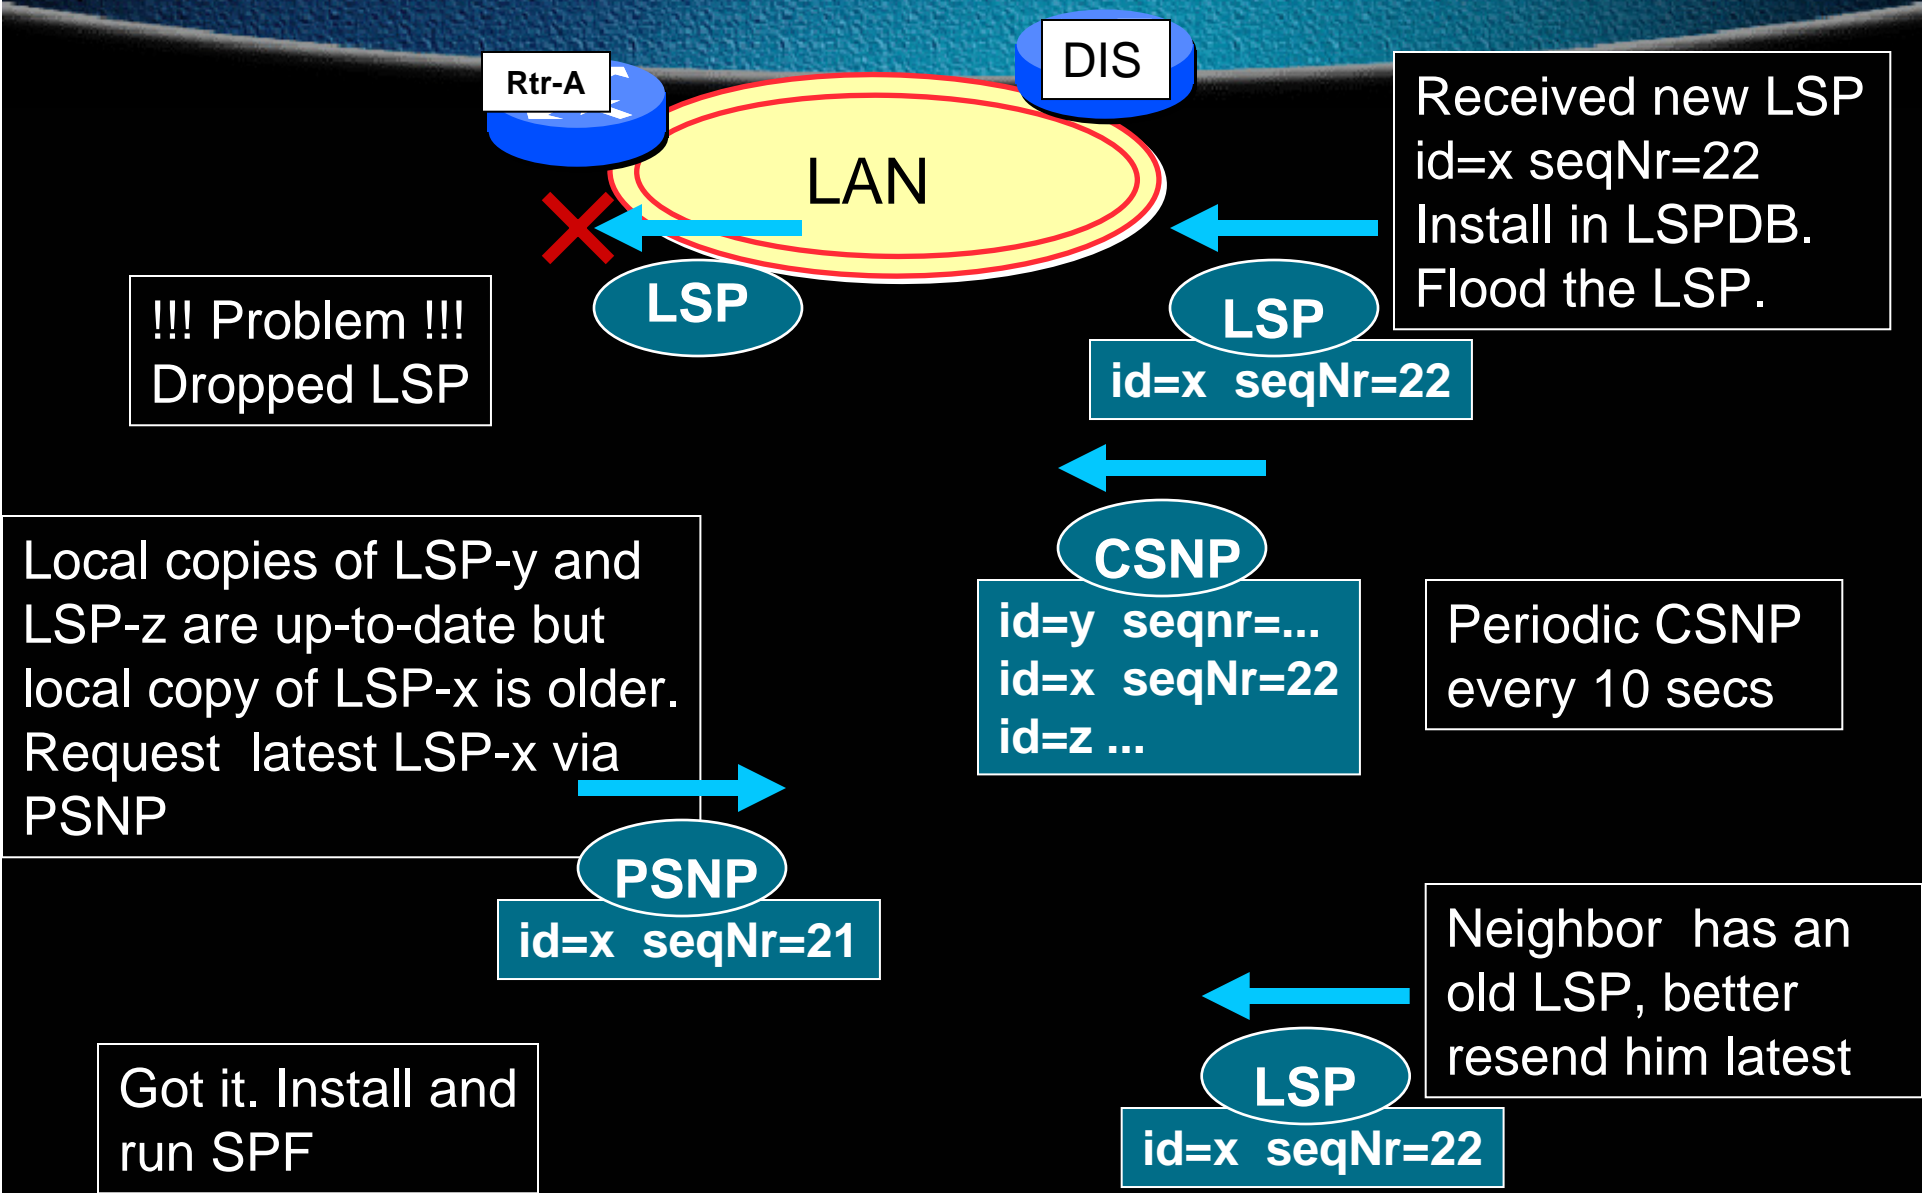
# The Designated Router

- **DIS is like the DR in OSPF**
- **DIS is only on LANs, not on p2p**
- **DIS has two tasks**
  - create/update pseudonode LSP**
  - conduct flooding over the LAN**
- **DIS sends persiodic CSNPs**
  - LSPid, SeqNr, Checksum, Lifetime of all LSPs present in the LSPDB**

# The Designated IS

- **No Backup DIS in ISIS**  
not necessary, no LSPDB resync
- **DIS is elected by priority and MAC**  
actually is “self-elected”
- **LAN circuitID shows who is DIS**  
use `show clns interface`

# Flooding on a LAN





# IP routing specifics

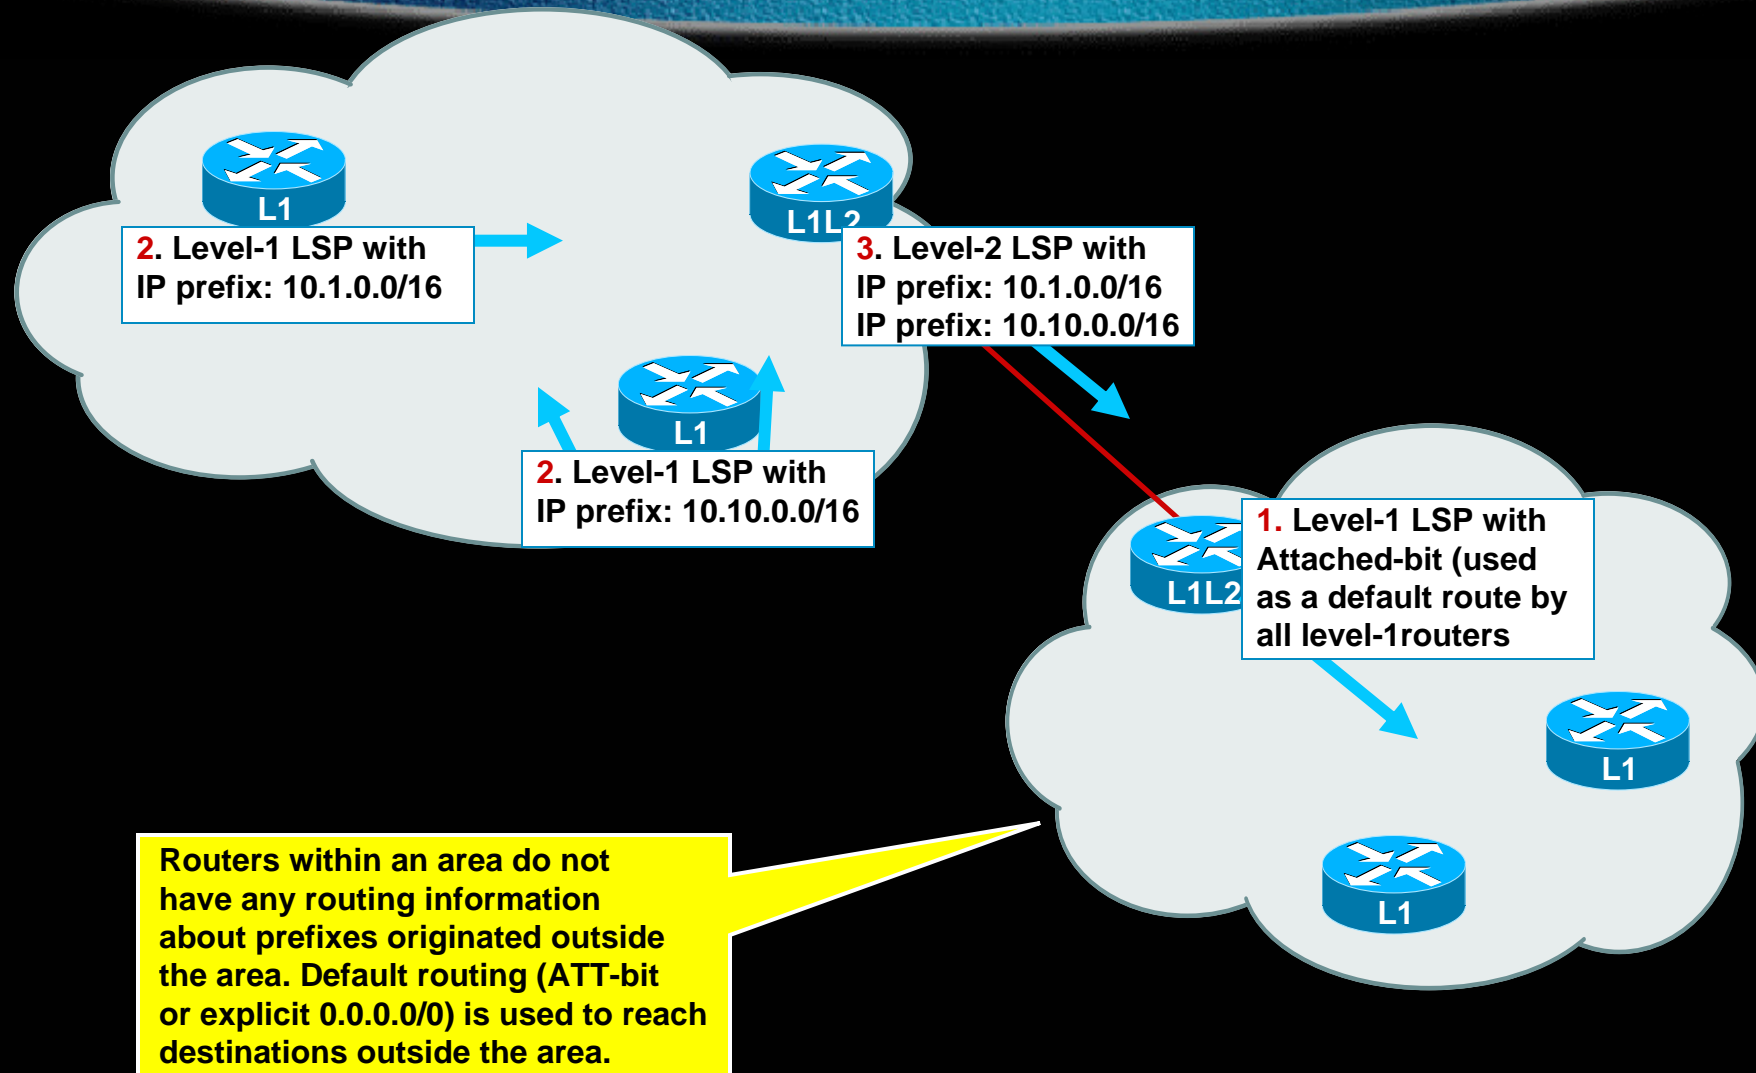
# L1 advertised into L2

- **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)**
- **Summarization possible**  
**At L1->L2 or when redistributing**

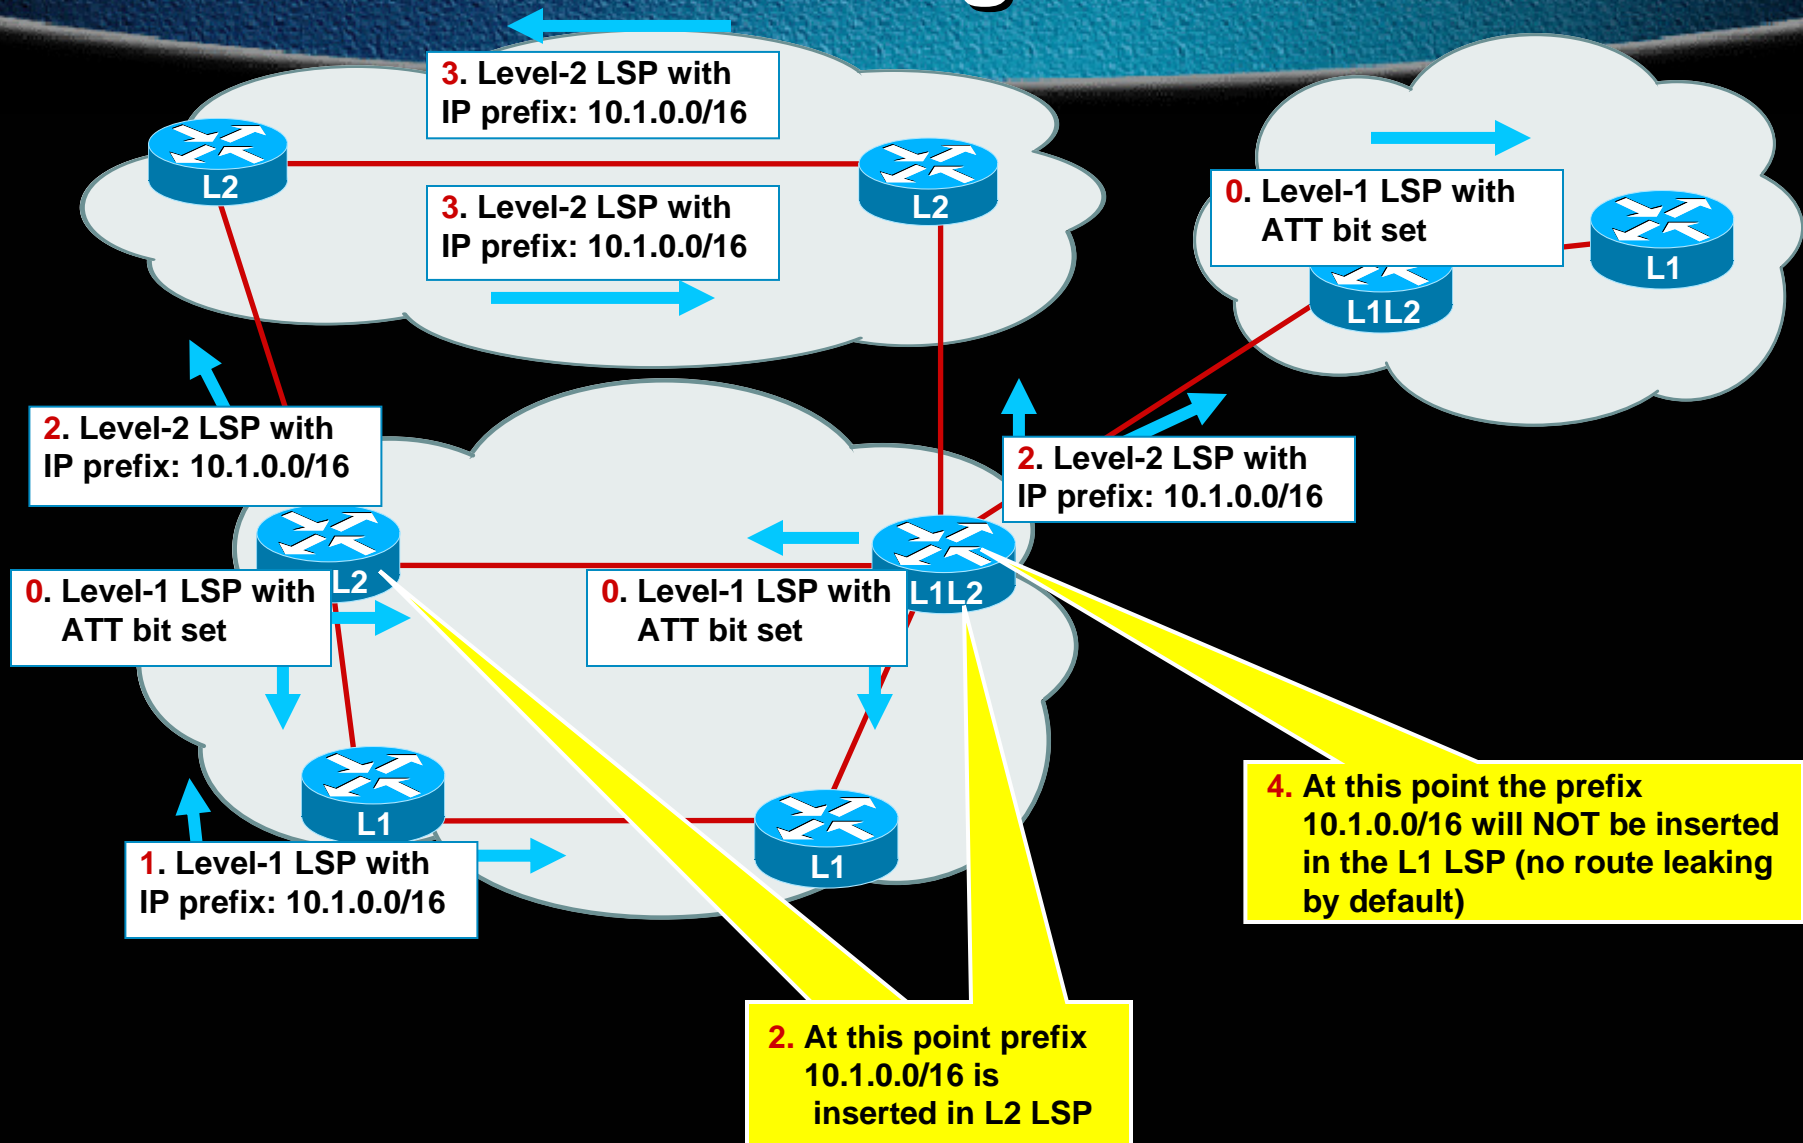
# L1 advertised into L2

- **Level-1 preferred over Level-2**  
In case of SAME routes (same prefix and same mask)
- **Internal equal to external route type**  
(TLV 128 versus TVL 130)
- **Internal over external metric-type**

# ISIS routing Levels



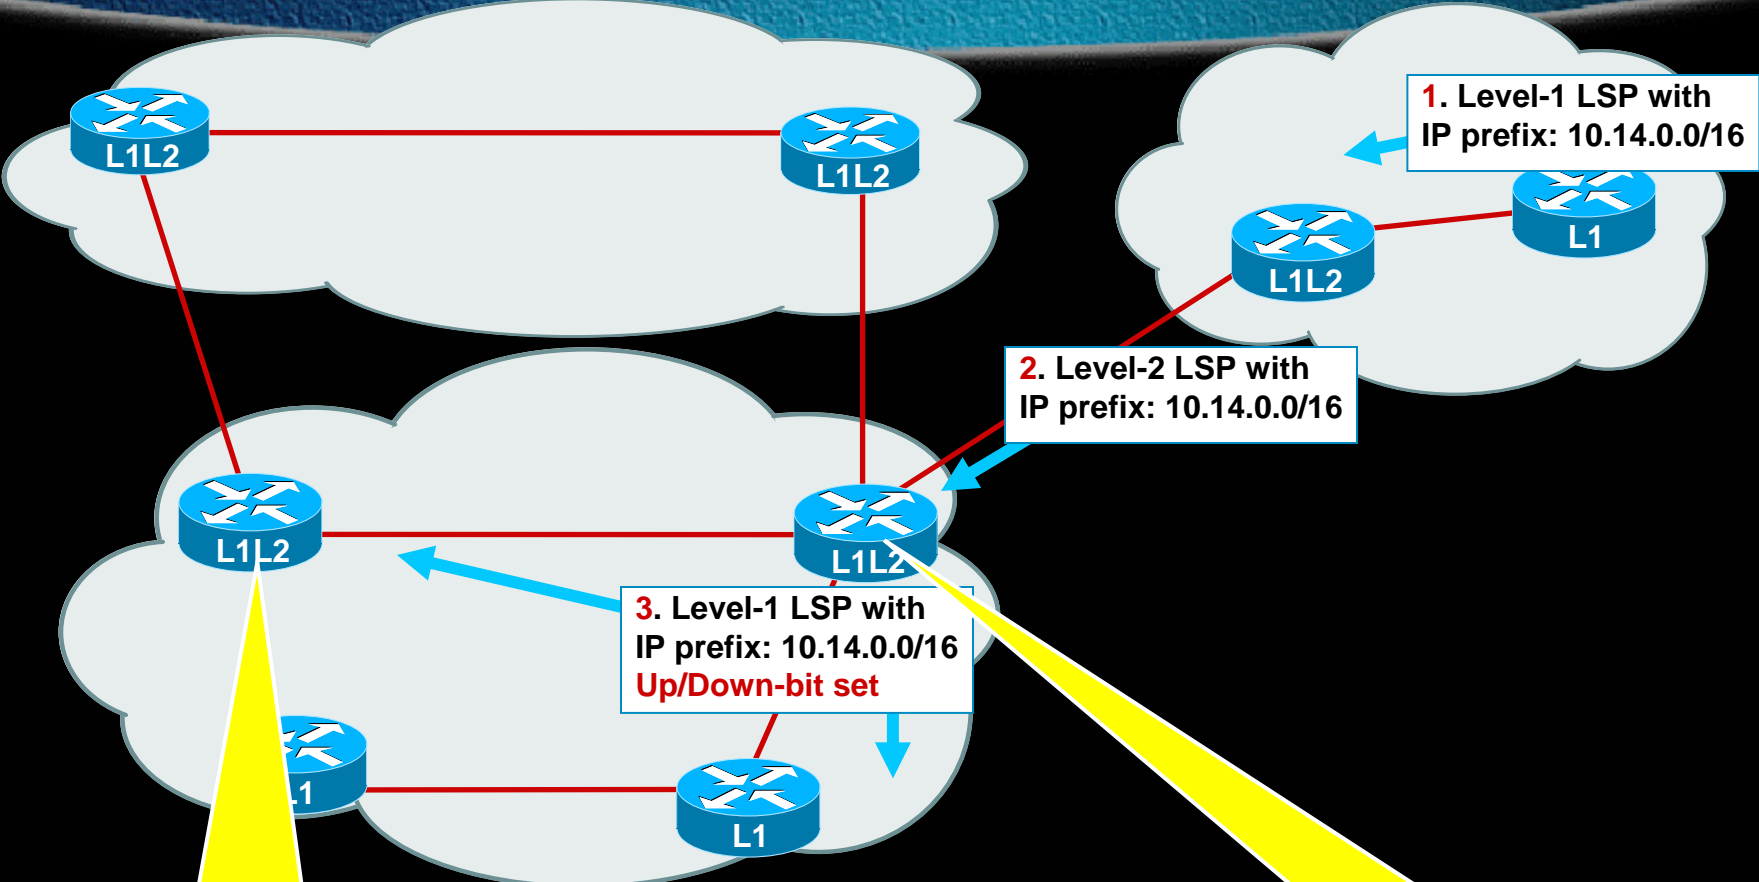
# ISIS routing Levels



# Route Leaking

- **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



1. Level-1 LSP with IP prefix: 10.14.0.0/16

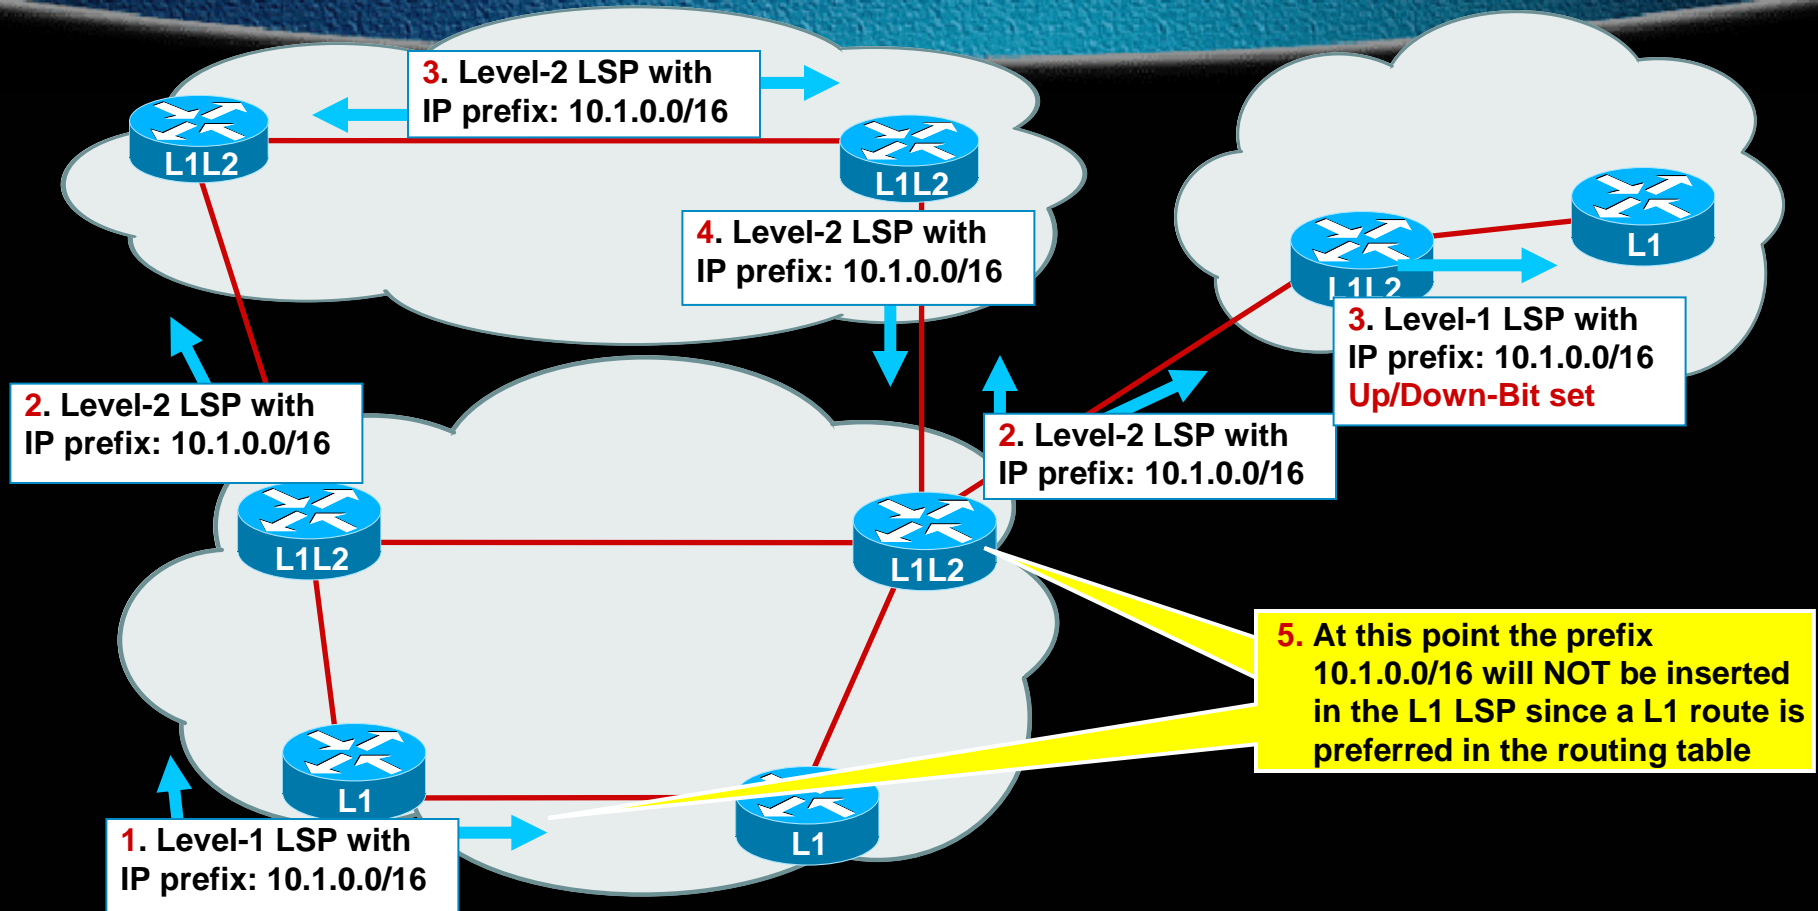
2. Level-2 LSP with IP prefix: 10.14.0.0/16

3. Level-1 LSP with IP prefix: 10.14.0.0/16 Up/Down-bit set

4. At this point prefix 10.14.0.0/16 will NOT be inserted in L2 LSP since it has the Down-bit set

3. At this point prefix 10.14.0.0/16 will be inserted in L1 LSP since route leaking is configured AND the prefix is present in the routing table as a L2 route

# Route Leaking



# Route Leaking

- **For IP only**
- **Prefixes MUST be present in the routing table as ISIS level-2 routes**

**Otherwise no leaking occurs**

**Same criteria than L1 to L2**

**Inter-area routing is done through the routing table**

# Route Leaking

- **Solution for several issues:**
- **optimal inter-area routing**
- **BGP shortest path to AS exit point**
- **MPLS-VPN (PEs loopback reachability)**

# Route Leaking

- **When leaking routes from L2 backbone into L1 areas a loop protection mechanism need to be used in order to prevent leaked routes to be re-injected into the backbone**

# Route Leaking

- **UP/Down bit**

Extended IP Reachability TLV (135) contains Up/Down bit

Described in draft-ietf-isis-traffic

- **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

- **Route leaking works also with**
  - IP Internal Reachability Information (TLV 128)**
  - IP External Reachability Information (TLV 130)**
  - Also known as: Old-style or Narrow-metric TLVs**

# Route Leaking

- **TVLs 128 and 130 have a metric field that consists of 4 TOS metrics**

**The first metric, the so-called "default metric", has the high-order bit reserved (bit 8) Routers must set this bit to zero on transmission, and ignore it on receipt**

- **The high-order bit in the default metric field in TLVs 128 and 130 becomes the Up/Down bit**

# Route Leaking

- **Recommendation:  
use wide Metric TLV (TLV 135)**
- **Configured with:**

***Router isis  
metric-style wide***

# Route Leaking

- **Route leaking is implemented in both 12.0S and 12.1**

**Cisco IOS 12.0S command**

*advertise ip L2-into-L1 <100-199>*

**Cisco IOS 12.1 command**

*redistribute isis ip level-2 into level-1 distribute-list <100-199>*

- **Both commands are supported**
- **12.0S command will be converted into 12.1 syntax**

# Summarization is possible .....

- From L1 areas into the L2 backbone,
- From L2 leaking down into L1 areas,
- When redistributing into L2 or L1

*router isis*

*summary address 192.1.0.0 255.255.0.0*

# Partial Route Calculation

- **Recompute all IP routes in a LSP**

**Only implemented for IP**

- **Done when SPT topology doesn't change**

**IP prefix added, removed or metric change**

**Inter-area, externals and passive interfaces**

# ISIS router configuration

*router isis*

*net 49.0001.0000.0000.000a.00*

*log-adjacency-changes*

*is-type level-2-only*

*passive-interface loopback0*

# ISIS interface configuration

*interface POS2/0*

*ip unnumbered loopback0*

*ip router isis*

*clns router isis*

*isis circuit-type level-2*

A man in a white shirt and dark tie is holding a long, curved cable against a blue background. The cable is arched over his head, and he is looking at it. The background has some faint, vertical lines.

# Network design issues

# System-ID addressing

- **System-ID is 6 bytes**
- **Convert your loopback IP addresses into CLNS System-Ids**
- **Or simply use:**

**0000.0000.0001**

**0000.0000.0002**

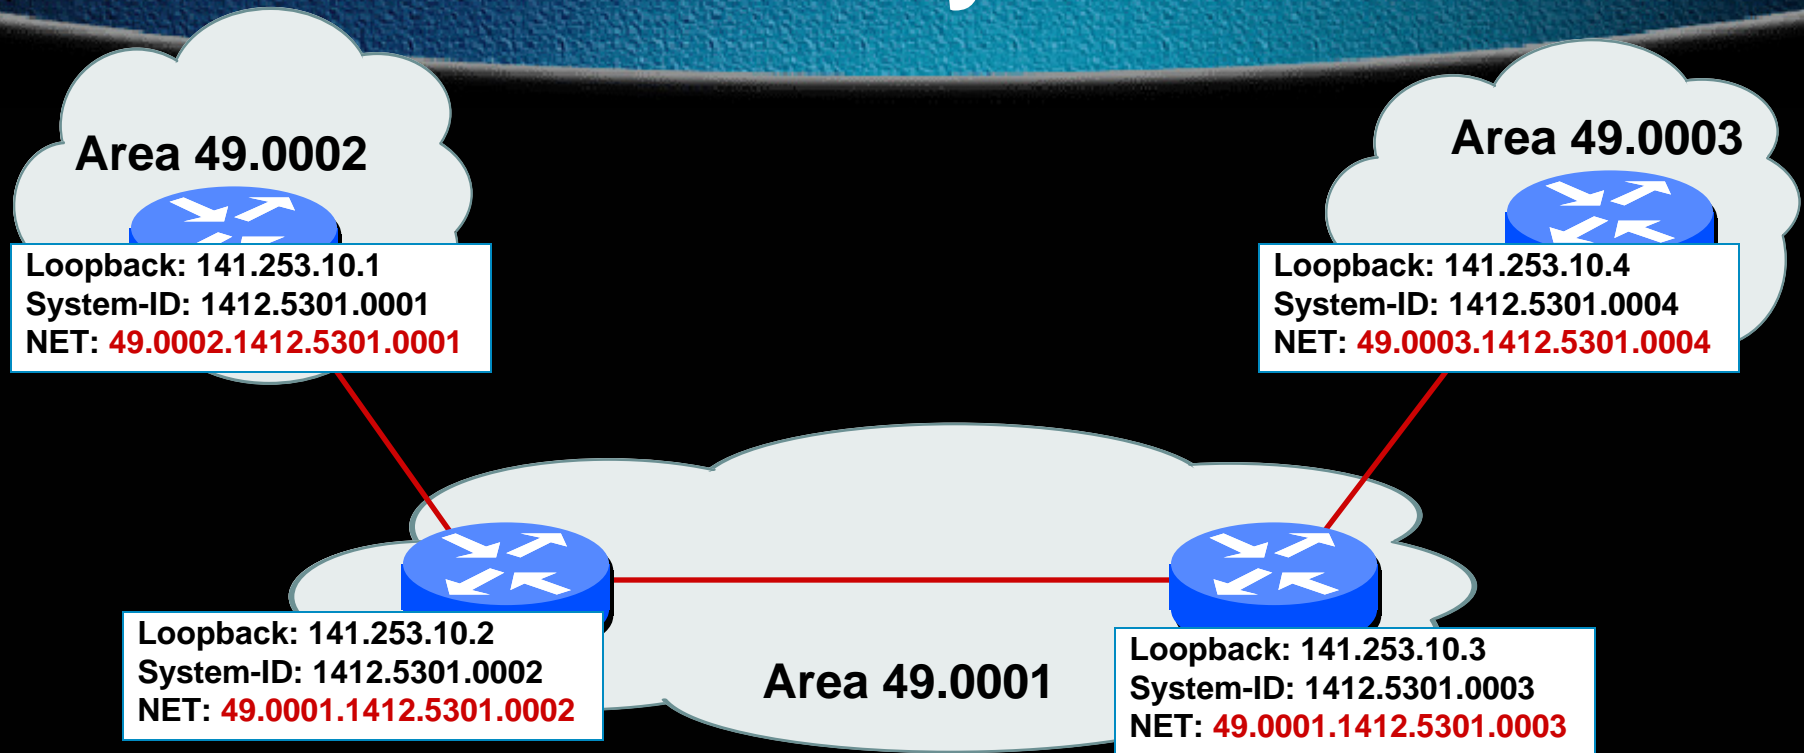
**0000.0000.0003**

**...**

# Area addresses

- **If PUBLIC CLNS network, request an official NSAP prefix**
- **Otherwise, use AFI 49**  
**Especially for IP networks**
- **AFI 49 is CLNS private address space**  
**like network 10.0.0.0 in IP**
- **Areas addresses:**  
**49.0001..., 49.0002..., 49.0003,... etc**

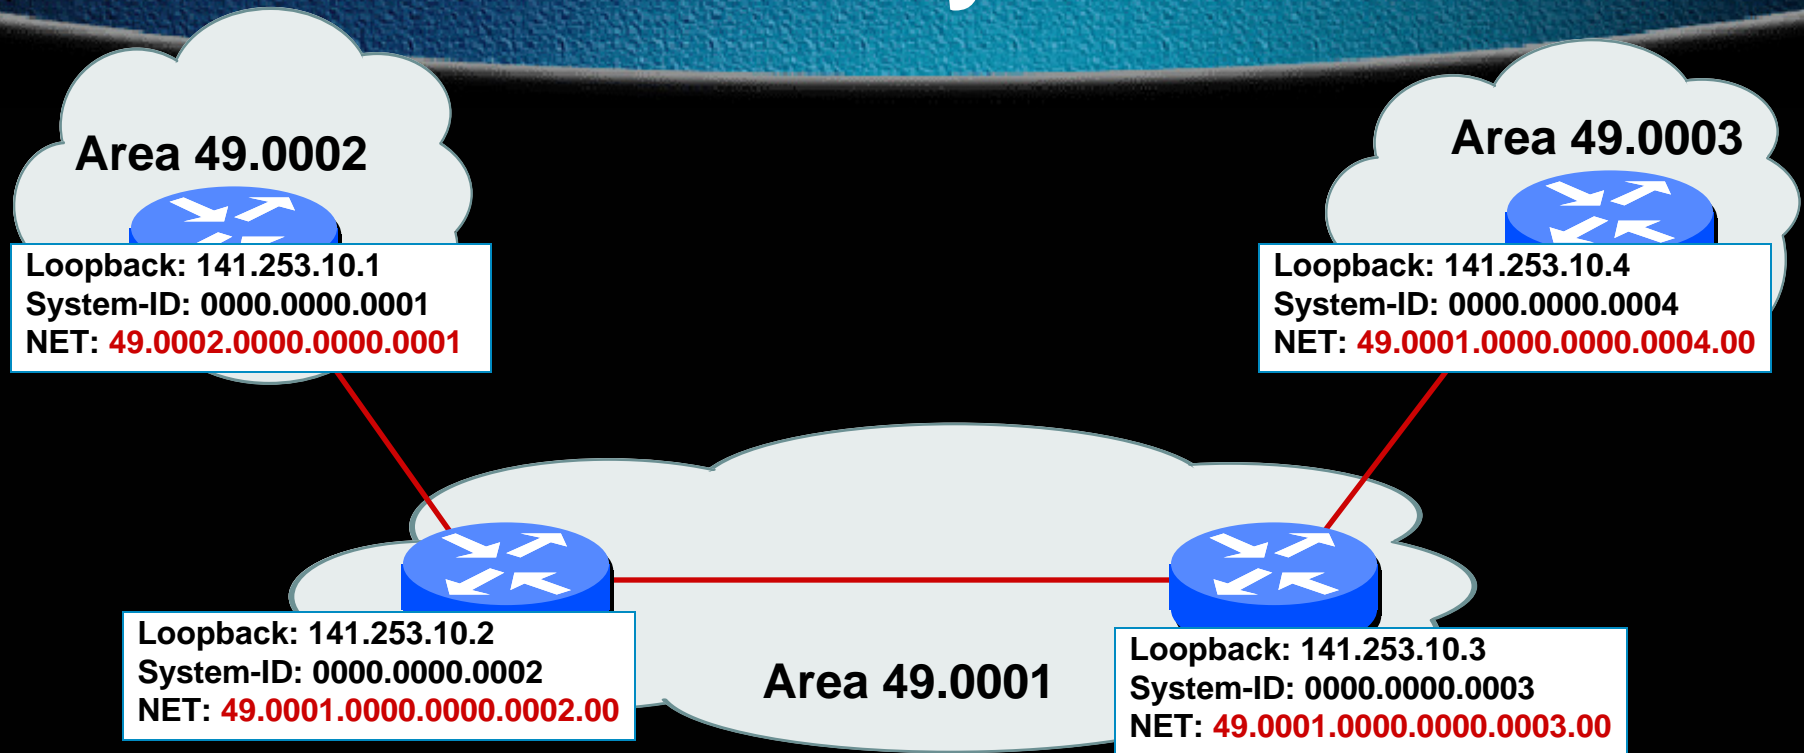
# Areas and System-IDs



- **System-Ids can easily be derived from Loopback addresses**

141.253.10.4 → 141.253.010.004 → 1412.5301.0004

# Areas and System-IDs



# Link-State protocols main components

- **SPF computation**
  - no big deal in most networks**
  - forwarding continues during SPF**
- **LSP flooding**
  - number of neighbors**
  - speed of links, nr of buffers**
  - size of the network**

# LSP flooding throttling

- **ISO 10589 states LSP flooding on a LAN should be limited to 30 LSPs/sec**  
**33 msec gap between consecutive LSPs**
- **IOS throttles over both LAN and p2p**
- **30 LSP/sec, 1492 byte/LSP = 360 Kbps**  
**slow speed links might get congested**  
**use “lsp-interval” command**

# LSP flooding throttling

- **Several interface configuration commands**

**isis lsp-interval 33**

delay (msecs) between LSP transmissions (flooding)

**isis retransmit-interval 5**

delay between retransmissions of the same LSP

**isis retransmit-throttle-interval 100**

delay between retransmitted LSPs

**isis mesh-group blocked**

Block LSP flooding on this interface

# LSP flooding over LANs

- LAN flooding usually doesn't encounter any problem
- No retransmission over LANs
- No ACKs on LANs

DIS only sends periodic CSNPs

- Reduce CSNP timer for faster convergence over a LAN

```
int ethernet 1/0
```

```
isis csnp-interval <0-65535>
```

# Advertising parallel p2p adjacencies only once

- When building an LSP all adjacencies are inserted from the adjacency DB
- Parallel p2p adjacencies may be advertised in the LSP more than once
- This does not matter for SPF
- Optimization: advertise parallel p2p adjacencies only once

# Advertising parallel p2p adjacencies only once

- **Several scaling advantages**

**LSPs will be smaller**

**SPF calculations will be more efficient**

**Flapping of one of a set of parallel links  
will be invisible to the rest of the network**

# Wider Metric Support

- **Default metric used is cost (default is set to 10 on all interfaces)**
- **The interface metric was limited to 6 bits wide (range 0 to 63)**
- **The total path metric was limited to 10 bits wide (maximum of 1023)**

# Wider Metric Support

- **The interface metric was increased to 24 bits wide (range 0 to 16777215)**
- **The total path metric was increased to 4261412864 (32-bits) from 1023**
- **Can configure the old or new metrics**
- **Default is old style metrics**

# Wider Metric Support

- **Configuration:**

**Rtr-A(config)#router isis**

**Rtr-A(config-router)#metric-style ?**

**narrow** Use old style of TLVs with narrow metric

**wide** Use new style of TLVs to carry wider metric

**Rtr-A(config-router)#metric-style wide ?**

**level-1** Level-1 only

**level-1-2** Level-1-2

**level-2** Level-2 only

# Design guidelines

## LSP Generation

- **LSP generation (lsp-gen-interval )**

**Control the “frequency” of LSP generation**

**Prevent from flapping links causing a lot of LSPs to be flooded throughout the network**

# Design guidelines SPF, PRC and LSP generation

- **IS-IS throttles its main events**
  - SPF/PRC computation, LSP generation**
- **Throttling slows down convergence**
- **Not throttling can cause melt-downs**
- **Find a compromise...**

# Design guidelines

## SPF, PRC and LSP generation

- **Initial wait before SPF/PRC or LSP generation**
  - 50 msec for LSP gen, 5.5 sec for SPF**
- **And a configurable minimum interval between consecutive events**

**By default is**

**5 seconds for LSP generation**  
**10 seconds for SPF calculation**

# Design guidelines

## SPF, PRC and LSP generation

- **The scope is to react fast to the first events but ,under constant churn, slow down to avoid to collapse**

# Design guidelines

## SPF, PRC and LSP generation

- **Extended syntax**

**spf-interval <a> [<b> <c>]**

**<a> seconds between SPF runs (seconds)**

**<b> milliseconds between first trigger and SPF**

**<c> milliseconds between first and second SPF**

- **Same syntax for**

**prc-interval**

**lsp-gen-interval**



# Design guidelines

## SPF, PRC and LSP generation

- **When the network calms down, and there were no triggers for 2 times the minimum interval (20 sec in this example), go back to fast behaviour (100 ms initial wait)**

# Design guidelines SPF and PRC

- **SPF timers (spf-interval & prc-interval)**

**Control the “frequency” of SPF runs**

**Prevent from collapsing the router if SPF has to be run constantly**

**Backoff can be configured however...  
danger !!!**

- **SPF Backoff is **\*\*NOT\*\*** a dampening mechanism**

# Design guidelines

## Lifetime

- Lifetime and refresh interval
- Set lifetime to its maximum value

*Router isis*

*max-lsp-lifetime 65535*

*lsp-refresh-interval 65400*

- Allow LSP refresh only after more than 18 hours
- Less flooding more stability

# Design guidelines

## Hello padding

- **IIHs are padded to full MTU size**
- **Benefit is early detection of problems**  
**transmission problems with large frames**  
**MTU mismatch**
- **On high speed interfaces MTU is 4K**  
**or higher and could be strain on huge**  
**buffers**
- **On low speed interfaces large hellos**  
**waste bandwidth**

# Design guidelines

## Hello padding

- Large hello packets waste bandwidth
- Can now be suppressed selectively

All interfaces

Per interface

*no hello padding*

# Design guidelines

## Default routing

- **ISIS uses the Attached-bit to discover the exit point of an area**
- **ATT bit is set into the L1 LSP by the router who is also a L2 router**
- **All L1 routers will use that L2 router as default exit point for the area**

# Design guidelines

## Default routing

- **ATT may not be optimal for IP networks**
- **Use explicit IP default route**  
*router isis*  
*default-information originate*
- **IP default route ALWAYS preferred over the ATT bit**

# Design guidelines

## Redistribution

- **The good practice is.....to avoid it !**
- **Otherwise can happen in any type of router (L1 or L2 or L1L2)**
- **Summarisation possible in the redistributing router**
- **Summarisation always possible in the L1L2 router for routes redistributed by another L1 router**
- **Different from OSPF**

# Route Leaking

- **New feature to enable redistributing level-2 IP routes into level-1 areas**

**Allow optimal area routing**

**IP only feature (CLNS still uses stub)**

- **Enables level-1-only routers to pick the best path to exit the area**
- **Enables shortest-exit and MED for BGP**
- **Enables MPLS-VPN (PE reachability) between areas**
- **Redistribution is controlled via distribute-lists**

# Route Leaking

- **Can use the new or old style metrics when leaking routes**

**Better use the new wider metrics to get real benefit from leaking routes into L1**

**Otherwise all backbone routes may be seen as advertised by L1L2 routers with a metric of 63 !**

**Route leaking benefit is lost....**

# Design guidelines

## Summarisation

- **Summarisation**
- **Used with area routing**
- **Used from level-1 area to level-2 backbone**
- **Used from level-2 backbone to level-1 area when using route leaking**
- **Always a good practice**

# Design guidelines Management & troubleshooting

- **Keep trace of**  
**SPF calculations**  
**Adjacencies states**
- **Show command**  
**show isis spf-log**
- **Configuration command**  
**log-adjacency-changes**

# Design guidelines

## Overload-bit

- **10589 defines for each LSP a special bit called the LSPDB Overload Bit**
- **While having problems, a router could set the OL bit, and other routers would route around it**
- **Connected IP prefixes still reachable**

**This may change in the future**

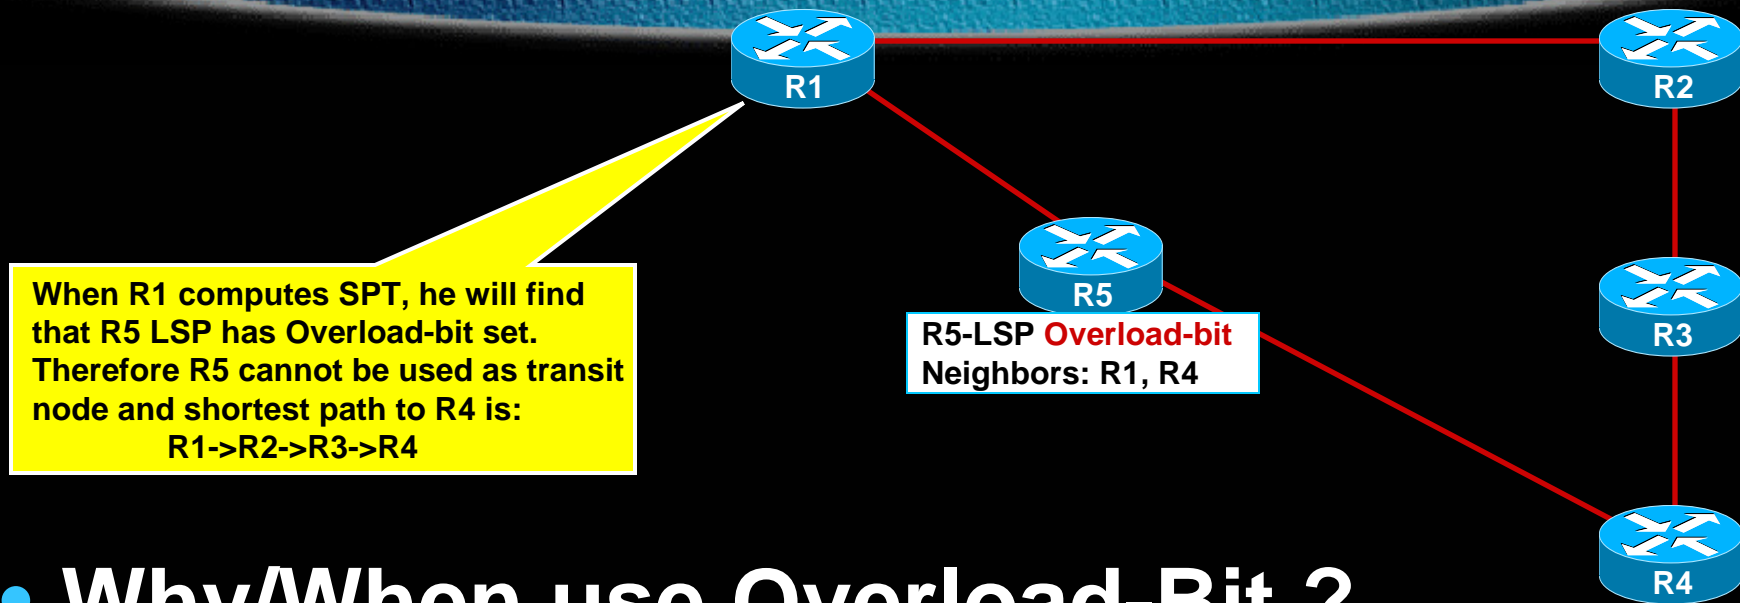
# Design guidelines

## Overload-bit

- **With IS-IS you can manually set the overload bit in the router's LSP**
- **This router will therefore never be used for transit during the path calculation, but it is still reachable**
- **Use for routers in the lab, routers aggregating management PVCs, etc**

# Design guidelines

## Overload-bit



- **Why/When use Overload-Bit ?**

**When the router is not ready to forward traffic for ALL destinations**

**Typically when ISIS is up but BGP (or even MPLS) not yet**

**When the router has other functions (Network Management)**

# Design guidelines

## Overload-bit

- **BGP will typically converge much slower than the IGP (a few minutes)**
- **During this time, other routers in the AS will use this new router for transit**
- **But if the new router does not have all BGP routes yet, it will drop traffic**
- **New router should first converge BGP before carrying traffic**

# Design guidelines

## Overload-bit

- **IS-IS can set the OL bit after each reboot, and allow BGP to converge before it advertises itself as transit by unsetting the OL bit**
- **Network admin needs to specify how long IS-IS should wait for BGP to converge**  
**typically 2 to 5 minutes**

# Design guidelines

## Overload-bit

- **BGP can tell IS-IS to unset the Overload-bit immediately**
- **Default BGP update delay is 2 min**
- **When BGP never informs ISIS, the Overload-bit will be cleared after 10 minutes**

# Design guidelines

## Overload-bit

***router isis***

***set-overload-bit***

***set-overload-bit on-startup <sec>***

***set-overload-bit on-startup wait-for-bgp***

***router bgp 100***

***bgp update-delay <sec>***

# Design guidelines

## Overload-bit

- **Overload-bit on-startup recommended in MPLS networks**
- **During boot-up a router may have all IGP routes but not all labels**
- **During this time it's better not to use the router as a transit point**

**router isis**

**set-overload-bit on-startup 120**

A man in a white shirt and dark tie is holding a long, curved object, possibly a cable or a pipe, against a blue background. The object is curved in a large arc, and the man is standing on a dark, curved surface. The overall scene is monochromatic with shades of blue and black.

# New features

# Implementation

- **Partial Route Calculation**
- **FastPSNP**
- **PRC on MPLS TE tunnel flap**
- **Parallal p2p adjacencies**

# Dynamic hostname to systemID resolution

- IS-IS routers and LSPs identified by a 6 byte systemID, taken from NSAP
- **Before 12.0(5)S and 12.0(5)T:**
- Name to NSAP mappings had to be manually configured:

```
Rtr-A(config)#clsns host <name> <nsap>
```

# Dynamic hostname to systemID resolution

- **New feature to automatically learn router names**
- **Symbolic name is put inside the LSP**
- **This feature is ON by default**

**can be disabled via:**

**router isis**

**no dynamic-hostname**

# Multi-Area Support

- **Allows multiple IS-IS processes to be configured on a single router**
- **Each process serves a different area**
- **One of the processes will be L1L2 to advertise all area addresses from all processes into L2**
- **Maximum number of configurable IS-IS processes on a single router is 29**

# Traffic Engineering Support

- **Support for MPLS-TE:**  
Traffic Engineering with Multi Protocol Label Switching
- **IS-IS allows MPLS-TE to flood resource, policy and reservation information about links inside LSPs**
- **New information carried in sub-TLVs**

# Faster link failure detection

- Holdtime is the number of seconds a neighbor must maintain an adjacency when it does not receive hellos
- Holdtime is the hello-interval multiplied by the hello-multiplier
- The holdtime is advertised inside IHLs  
no need for matching configurations

# Faster holdtime

- Hold-time can be set to 1 second

*Interface POS2/0*

*isis hello-interval minimal*

- By default hello-multiplier is 3

**Hello packets sent every 333 msec**

# Fast Hellos

- **Advantages**

**reduced link failure detection time**

- **Disadvantages**

**increased BW/buffer/CPU usage can cause missed hellos. potential increased adjacency flapping can cause instability**

**use no isis hello padding feature to reduce BW and buffer usage**

*interface Pos 2/0  
no isis hello padding*

# Fast Hellos

- **Configuration:**

**Interface configuration mode:**

**Rtr-A(config)#int POS2/0**

**Rtr-A(config-if)#isis hello-interval minimal**

**Rtr-A(config-if)#isis hello-multiplier 5**

- **Advertised holdtime will now be 1 second, hello-interval will be 200 ms**



# Future enhancements

# Future Enhancements

- **Route TAGS**

**Possible new sub-TLV for tagging**

**Similar to BGP communities**

**Allows flexible filtering at area boundaries**

# Potential enhancements

- **Flooding optimization over parallel adjacencies**
- **Extensions for DPT (SRP)**
- **ISIS for IPv6**
- **Incremental SPF ?**
- **More level hierarchy ?**



# The Integrated ISIS routing protocol

Session Number

2222



# Please Complete Your Evaluation Form

Session Number

**2222**

# CISCO SYSTEMS



EMPOWERING THE  
INTERNET GENERATION<sup>SM</sup>