

Agenda

- Differences between IS-IS and OSPF
- SPF Algorithm
- CLNS Addressing
- IS-IS Protocol Concepts
- IS-IS LSDB

Integrated IS-IS and OSPF

**Some subtle (and not so subtle
differences...)**

Similarities

- Link State Protocols
- Link Types
- Hello Protocol
- Two Level Hierarchy
- Designated Router on LAN
- Interoperable
- Widely Deployed
- Simple Password and MD5 Authentication

Differences

Encapsulation

- **OSPF**
 - Runs over IP
 - Partitioning (Virtual Links)
 - Spoofing and attacks
 - ATM VCMux Encap
- **IS-IS**
 - Runs over L2
 - Tunneling for Partition Repair (Difficult!)
 - Difficult to Spoof
 - ATM SNAP Encap

Differences

Protocols Supported

- **OSPF**
 - IP
- **IS-IS**
 - IP
 - ISO-CLNP
 - ...

Differences

Link Types

- **OSPF**
 - Point-to-Point
 - Broadcast
 - NBMA
- **IS-IS**
 - Point-to-Point
 - Broadcast

Differences

Router Identification

- **OSPF**
 - Router-id (e.g. 192.168.1.1)
 - Highest IP Address on router
 - “router-id” statement
- **IS-IS**
 - NSAP / NET (e.g. 49.0001.1921.6800.1001.00)
 - Comprise of Area ID, System ID, NSEL
 - Explicitly Configured

Differences

Adjacencies (Broadcast)

- **OSPF**
 - Designated Router (DR)
 - Backup Designated Router (BDR)
 - Interface Priority = 1
 - Non-Preemptive
- **IS-IS**
 - Designated Integrated System (DIS)
 - Interface Priority = 64
 - Pre-emptive

Differences

Hello Protocol

- **OSPF**
 - Hello-Interval = 10 s
 - Dead-Interval = 40 s
 - Timers & MTU must match!
- **IS-IS**
 - Hello Padding
 - Hello-Interval = 10 s (DIS = 3.3 s)
 - Hold-Time = 30 s

Differences

Individual Link Cost (Metric)

- **OSPF**
 - 1-65535
 - Link Bandwidth / Ref BW
- **IS-IS**
 - 1-63 (Default Metric)
 - 1-16777214 (Metric-Style Wide)
 - Fixed cost of 10

Differences

Packet Encoding


- **OSPF**
 - 32-bit Aligned
 - Only LSAs are extensible
 - Only known LSAs are flooded
 - Opaque LSAs
- **IS-IS**
 - No Alignment
 - TLVs - Extensible
 - Unknown TLVs are flooded

Differences


Areas

- **OSPF**
 - Boundary on Routers
 - ABR for Multiple areas
- **IS-IS**
 - Boundary on Links
 - Router in single area + Backbone

Differences - Areas (OSPF)



Differences - Areas (IS-IS)



Differences

Database


- **OSPF**
 - Link State Advertisement (LSA)
 - Age incremented till *max-age* in LSDB
 - Refresh every 30 min
- **IS-IS**
 - Link State Packet (LSP)
 - Originator tags LSAs with lifetime (0 to 65535s)
 - Decrement till expiry
 - Refresh every 15 min

Shortest Path First Algorithm

Shortest Path First Algorithm

- Also known as “**Dijkstra**”
- Computes the shortest path to all reachable nodes
- We maintain three lists
 - **UNKNOWN** list: All nodes start on this list
 - **TENTative** list: All nodes to which a path has been found which is not yet known to be **minimum** cost
Also called the **candidate** list
 - **PATHS** list: All nodes to which we have calculated final paths with definitive cost
Also called the **known** list

Shortest Path First Algorithm



Shortest Path First Algorithm

- We execute **N** steps
 - Typically **N** is the number of nodes in the network; during each step we find the path(s) to one node from root
- At each step:
 - Find the node amongst all nodes on **TENT** that has the lowest cost, and move it from **TENT** into **PATHS**
 - Find all neighbours reachable from that node and move them into **TENT**
 - Find all prefixes advertised by the node that was moved to **PATHS** and install them in the **RIB**

Shortest Path First Algorithm: IP Prefixes

- During SPF, when a node is moved into PATHS, all IP prefixes advertised by that node are inserted into the routing table
- In IS-IS, IP prefixes are **leaves** of the shortest path tree
 - We **don't use IP prefixes** to build the SPT
 - Routers are identified through **CLNS System-IDs**

Shortest Path First Algorithm:

SPF and PRC

- Basic Implementation
 - **Any change (link, node, leave)**
 - → Recompute the whole SPT and the whole RIB
- Decouple SPT and RIB
 - **If any topology change (node, link)**
 - → Recompute SPT and the RIB
 - **If only an IP prefix changes**
 - → Keep the SPT, just update the RIB for the nodes whose prefixes have changed
- More on this later...

Called
“SPF”

Called
“PRC”

CLNS Addressing

CLNS Addressing

NSAPs and Addressing

- ISO/IEC 10589 distinguishes only 3 fields in the NSAP address format



- CLNS addressing comprise of 3 main parts:
 - Area Address - variable
 - System ID - 6 bytes
 - NSEL(ector) - value is 0x00 on a router
- NSAP of a router is called a NET

CLNS Addressing 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 (selector=0)
- **Multiple NETs are like secondary IP addresses; only use them when merging or splitting areas**

CLNS Addressing

AFI Values

Address Domain	AFI Value
X.121	37
ISO DCC	39
ISO 6523	47
Local	49

- X.121 - Int'l plan for public data networks
- ISO DCC - Data country code
- ISO 6523 ICD - Telex
- Local - For local use within network domain only

CLNS Addressing

Requirements and caveats

- **At least one NSAP is required per node**
- **All routers in the same area must have a common Area ID**
- **Each node in an area must have a unique System ID**
- **All level 2 routers in a domain must have unique System IDs relative to each other**
- **All systems in given domain must have System IDs of the same length in their NSAP addresses**

CLNS Addressing

Requirements and caveats

- **Multiple NSAPs allowed on Cisco routers for merging, splitting or renumbering**
- **All NSAPs on the same router must have the same system ID.**
- **The maximum size of an NSAP is 20 bytes**
- **Minimum of 8 bytes allowed on Cisco routers**
 - 1 byte for area, 6 bytes for system ID and 1 byte for N-selector
 - AFI prefix recommended to make minimum of 9 bytes

CLNS Addressing NSAP Examples

- 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 NSAP Examples



CLNS Addressing

Addressing convention

- How do most ISPs define System IDs? E.g.

The **LOOPBACK** IP Address: **192.168.3.25**

The **AREA** the Router Under Is: **49.0001**

Conversion Process:

192.168.3.25

192.168.003.025

1921.6800.3025

49.0001.1921.6800.3025

IS-IS Protocol Concepts

IS-IS Protocol Concepts


- **Nodes**
- **Links**
- **Areas and Adjacencies**
- **Level-1 and Level-2 routing**

IS-IS Protocol Concepts: Nodes

Network Nodes:

- **Hosts**
- **Level-1 Routers**
- **Level-2 Routers**
- **Level-1 and Level-2 Pseudonodes on broadcast links only**

IS-IS Protocol Concepts: Nodes




- Broadcast link represented as Pseudonode (PSN)
- PSN role played by the Designated Router (DIS)
- DIS election is preemptive (interface priority/highest MAC)
- IS-IS has only one DIS. DIS/PSN functionality supports database synchronization between routers on a broadcast type link.


IS-IS Protocol Concepts: Links

- **Network/Link Types:**
 - **Point-to-Point Links** (PoS, p2p FR, ATM)
 - **Broadcast Links** (Multi-point/multi-access, e.g. Ethernet, mp FR, mp ATM or RPR)
- **Non-broadcast multi-access treated as broadcast or point-to-point by configuration**

IS-IS Protocol Concepts: Areas



IS-IS Protocol Concepts: Areas



- IS-IS supports 2-level routing hierarchy
- Routing domain is carved into areas. **Routing in an area is level-1. Routing between areas is level-2**
- All ISO 10589/RFC1195 areas are stubs

IS-IS Protocol Concepts

- **IS-IS Packet Types**
 - **IS-IS Hello Packets (IIH)**
 - Level 1 LAN IS-IS Hello
 - Level 2 LAN IS-IS Hello
 - Point-to-point Hello
 - **Link State Packets (LSP)**
 - Level 1 and Level 2
 - **Complete Sequence Number packets (CSNP)**
 - Level 1 and Level 2
 - **Partial Sequence Number Packets (PSNP)**
 - Level 1 and Level 2

IS-IS Protocol Concepts: Adjacency

- **IS-IS Hellos (IIH)** advertised for detecting neighbor and forming network layer adjacencies
- 3 types of IIHs:
 - **Point-to-point**
 - **LAN Level-1**
 - **LAN Level-2**

IS-IS Protocol Concepts: L1/L2

- IS-IS has a 2 layer hierarchy
 - The backbone (Level 2)
 - The areas (Level 1)
- An IS can be
 - Level 1 router (intra-area routing)
 - Level 2 router (inter-area routing)
 - Level 1-2 router (intra and inter-area routing)
- For each level (1 and 2) a DIS will be elected on LANs

IS-IS Protocol Concepts: L1/L2

Level 1 router

- Has neighbors only on the same area
- Has the Level 1 LSDB with all routing information for the area
- Use the closest Level 2 router to exit the area
- This may result in sub-optimal routing

Level 2 router


- May have neighbors in other areas
- Has a Level 2 LSDB with all information about inter-area routing

IS-IS Protocol Concepts: L1/L2

Level 1-2 router

- May have neighbors on any area
- Has two LSDBs:
 - Level 1 for the intra-area routing
 - Level 2 for the inter-area routing
- If the router has adjacencies to other areas, it will inform the Level 1 routers (intra-area) it is a potential exit point for the area

IS-IS Protocol Concepts: L1/L2



- Backbone **must be** L2 contiguous


IS-IS Link State Database

IS-IS Link State Database

- **Link State Packets (LSP)**
- **Sequence Number Packets**
- **IS-IS Metrics**
- **Flooding**
- **Database Parameters, Flags and Timers**

IS-IS LS Database

Link State Packets



IS-IS LS Database

IS-IS Packet Format

- IS-IS Packets are made up of:
- **A Fixed Header**
 - Contains generic packet information and other specific information about the packet
- **Type, Length, Value (TLV) Fields**
 - TLVs are blocks of specific routing-related information in IS-IS packets

IS-IS LS Database

Generic Packet Format

				No. of Octets
Intra-domain Routing Protocol Disc.				1
Length Indicator				1
Version/Protocol ID Extension				1
ID Length				1
R	R	R	PDU Type	1
Version				1
Reserved				1
Maximum Area Addresses				1
Packet-Specific Header Fields				
TLV Fields				

IS-IS LS Database

LSP Format

				Octets
Intradomain Routing Proto Discriminator				1
Lenth Indicator				1
Version/Protocol ID Extension				1
ID Length				1
R	R	R	PDU Type	1
Version				1
Reserved				1
Maximum Area Addresses				1
PDU Length				2
Remaining Lifetime				2
LSP ID				ID Length + 2
Sequence Number				4
Checksum				2
				1
P	ATT	LSPDBOL	IS Type	Variable
TYPE LENGTH VALUE FIELDS				

IS-IS LS Database

LS Packet ID (LSID)

00c0.0040.abcd.02-01



- **Sys ID** - System ID of IS or DIS in case of PSN LSP.
- **PSN ID** - Zero for Non-PSN LSP and Non-zero for PSN LSP
- **LSP Number** - Fragmentation number

IS-IS LS Database

Complete Sequence Number Packets (CSNP)

- **Describes all known LSPs in the LS Database**
- **Standard IS-IS packet format**
 - Fixed Header and
 - TLV Fields
- **Each TLV includes the following info about each known LSP:**
 - LSP ID
 - Remaining Lifetime
 - Sequence number
 - Checksum

IS-IS LS Database

Partial Sequence Number Packets (PSNP)

- **Similarly constructed as CSNPs**
- **Summaries of only specific LSPs**
- **PSNPs are used in 2 ways:**
 - To acknowledge receipt of LSP on pt-to-pt links
 - To request transmission of the latest version of an LSP

IS-IS LS Database

Level-1 TLVs

TLV Name	Type	Origin
Area Address	1	ISO 10589
Intermediate System Neighbors	2	ISO 10589
End System Neighbors	3	ISO 10589
Authentication information	10	ISO 10589
IP Internal Reachability Information	128	RFC 1195
Protocols Supported	129	RFC 1195
IP Interface Address	132	RFC 1195

IS-IS LS Database

Level-2 TLVs

TLV Name	Type	Origin
Area Address	1	ISO 10589
Intermediate System Neighbors	2	ISO 10589
Partition Designated Level-2 IS	4	ISO 10589
Prefix Neighbors	5	ISO 10589
Authentication information	10	ISO 10589
IP Internal Reachability Information	12	RFC 1195
Protocols Supported	129	RFC 1195
IP External Reachability Information	130	RFC 1195
Inter-domain Routing Protocol Information	131	RFC 1195
IP Interface Address	132	RFC 1195

IS-IS LS Database

New TLVs

TLV Name	Type	Comments
Extended IS Reachability Information	22	TE Extension. Replaces TLV 2
Router-Id	134	TE Extension
Extended IP Reachability Information	135	TE extension used in place of TLV 128 or 130
Dynamic Hostname Information	137	For dynamic distribution of Hostname to NET mapping via LSP Flooding
Point-to-Point Adjacency State	240	Reliable Pt-to-pt adjacency formation

IS-IS LS Database

Old IS-IS Metrics

- **ISO 10589 specifies 4 metric types**
 - **Default - Supported by all routers.**
 - **Delay - measures transit delay**
 - **Expense - Measures the monetary cost of link utilization**
 - **Error - measures error probability**
- **Default metric type must be supported by all implementations**
- **Other metric types specified for QoS routing. Not supported in Cisco IOS and most other vendor implementations**

IS-IS LS Database

Old IS-IS Metrics

- **Metric Information is carried in the following TLVs:**
 - IS Neighbors
 - ES Neighbors
 - Prefix Neighbors
 - IP Internal Reachability
 - IP External Reachability
 - Extended IP Reachability
 - Extended IS Reachability

IS-IS LS Database

Old IS-IS Metrics

0	I/E	Default Metric (6 bits)	1
S	I/	Delay Metric	1
S	E/	Expense Metric	1
S	I/E	Error Metric	1
IP Address			4
Subnet Mask			4

- Maximum metric per interface is 63
- Maximum path metric is 1023
- There is no automatic assignment based on interface bandwidth
- Cisco uses default of 10 on all interfaces regardless of bandwidth

IS-IS LS Database

Flooding and LSDB Synchronization

- Acknowledgement-based reliable mechanism on pt-to-point links
- Flooding on broadcast links depends on periodic advertisements of CSNPs for reliable flooding and synchronization

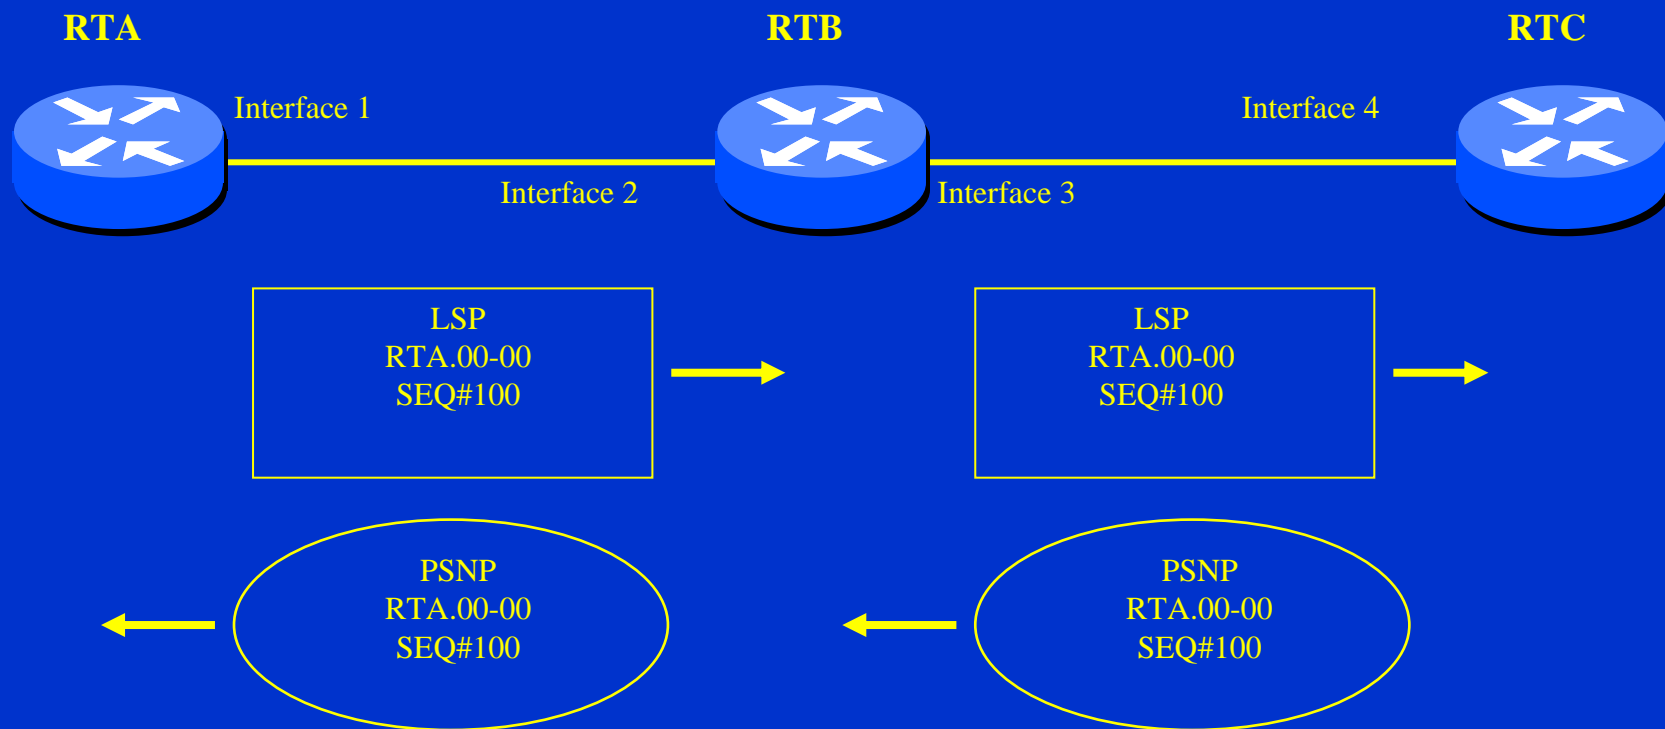
IS-IS LS Database

SRM and SSN Flags

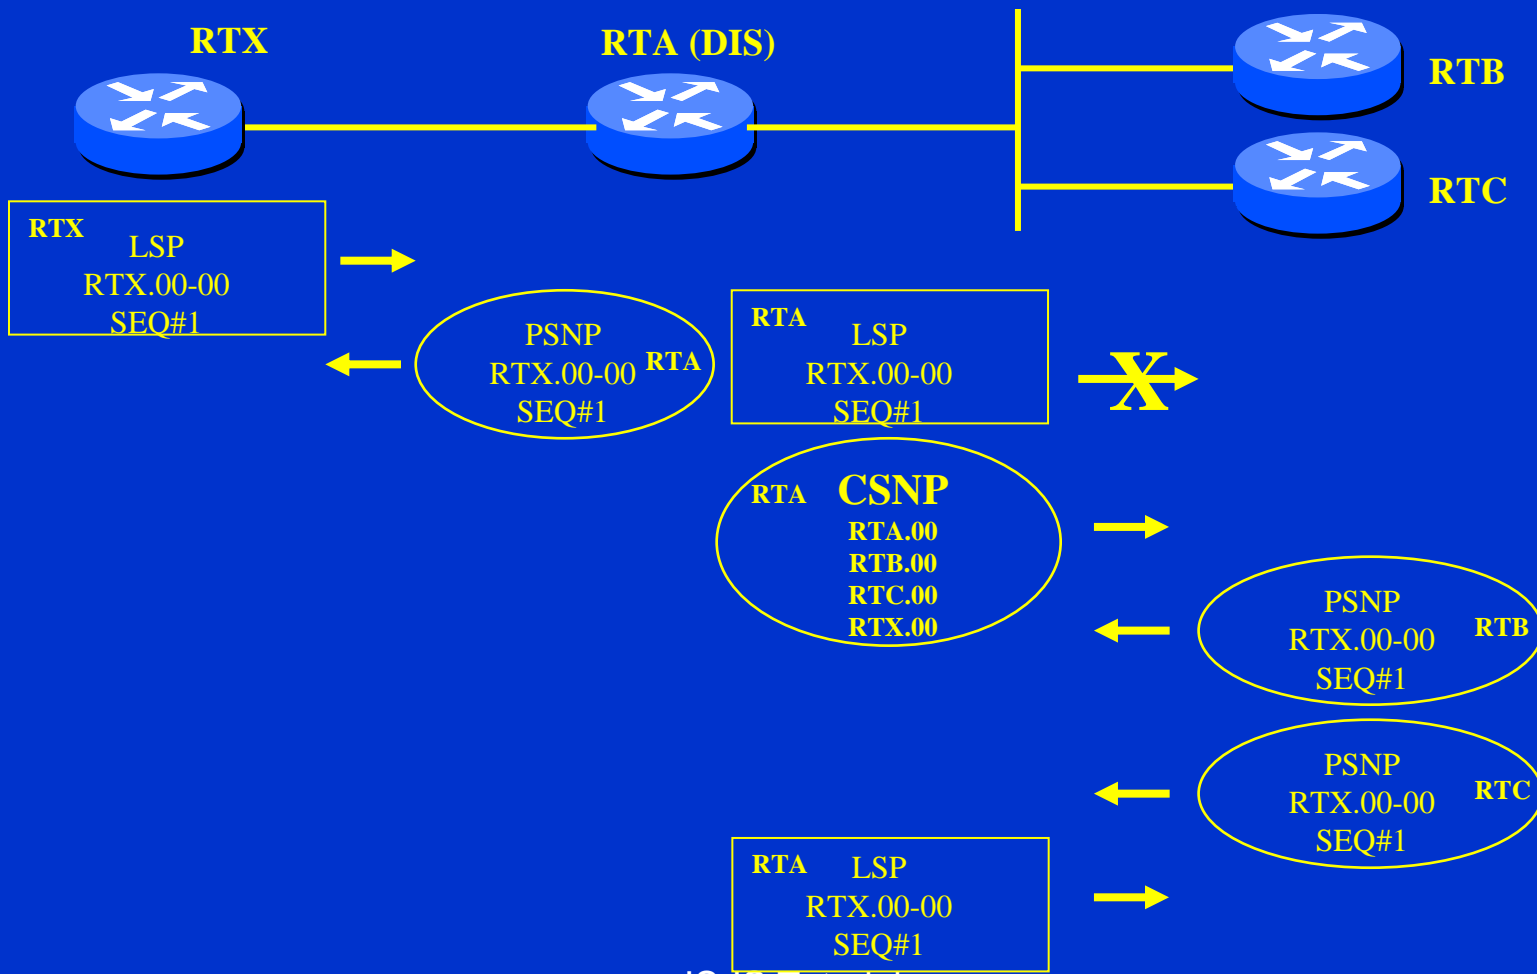
- **Send Routing Message flag (SRMflag)** and **Send Sequence Number flag (SSNflag)** are used by the Update Process to control delivery and acknowledgement of LSPs
- **SRMFlag** - If set for a LSP on a specific link, indicates the LSP should be flooded on that link
- **SSNFlag** - If set for a LSP on a link, means send a PSNP over that link to acknowledge receipt of LSP

IS-IS LS Database

Flooding on Point-to-Point Links



IS-IS LS Database Flooding on Broadcast Links



IS-IS LS Database

IS-IS Database Timers

Timer	Default Value	Cisco IOS Command
Maxage	1200s	isis max-lsp-interval
LSP Refresh Interval	900s	isis refresh-interval
LSP Transmission Interval	33ms	isis lsp-interval
LSP Retransmit Interval	5s	isis retransmit-interval
CSNP Interval	10s	isis csnp-interval