

- IS-IS -

IS-IS Fundamentals

IS-IS (Intermediate System -to- Intermediate System) is a standardized link-state protocol that was developed to be the definitive routing protocol for the **OSI (Open Systems Interconnect) Model**, which was developed by **ISO (International Standards Organization)**. IS-IS shares many similarities to OSPF. Though it was designed as an interior gateway protocol (IGP), IS-IS is predominantly used by ISPs, due to its scalability.

IS-IS adheres to the following Link State characteristics:

- IS-IS allows for a hierarchical network design using **Areas**.
- IS-IS will form **neighbor** relationships with adjacent routers of the same IS-IS type.
- Instead of advertising the distance to connected networks, IS-IS advertises the status of directly connected **“links”** in the form of **Link-State Packets (LSPs)**. IS-IS will only send out updates when there is a change to one of its links, and will *only* send the change in the update.
- IS-IS uses the **Dijkstra Shortest Path First** algorithm to determine the shortest path.
- IS-IS is a classless protocol, and thus supports VLSMs.

Other characteristics of IS-IS include:

- IS-IS was originally developed to route the ISO address space, and thus is *not* limited to IP routing.
- IS-IS routes have an administrative distance is **115**.
- IS-IS uses an arbitrary **cost** for its metric. IS-IS additionally has three optional metrics: **delay**, **expense**, and **error**. Cisco does not support these optional metrics.
- IS-IS has no hop-count limit.

The IS-IS process builds and maintains three separate tables:

- A **neighbor table** – contains a list of all neighboring routers.
- A **topology table** – contains a list of *all* possible routes to all known networks within an area.
- A **routing table** – contains the *best* route for each known network.

IS-IS is only available on **enterprise** versions of the Cisco IOS.

* * *

All original material copyright © 2007 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

IS-IS Protocols and Addressing

IS-IS consists of three sub-protocols that work in tandem to achieve end-to-end routing which ISO defined as **Connectionless Network Service (CLNS)**:

- **CLNP (Connectionless Network Protocol)** – serves as the Layer-3 protocol for IS-IS (and was developed by ISO).
- **ES-IS (End System -to- Intermediate System)** – used to route between *hosts* and *routers*.
- **IS-IS (Intermediate System -to- Intermediate System)** – used to route between *routers*.

IS-IS was originally developed to route ISO CLNP addresses (outlined in RFC 1142). However, CLNP addressing never became prominently used. Thus, IS-IS was modified to additionally support IP routing, and became **Integrated (or Dual) IS-IS** (outlined in RFC 1195).

The IS-IS CLNP address is **hexadecimal** and of **variable length**, and can range from **64** to **160 bits** in length. The CLNP address contains three “sections,” including:

- **Area** field – (variable length)
- **ID** field – (from 8 to 64 bits, though usually 48 bits)
- **Selector (SEL)** field - (8 bits)

Thus, the CLNP address identifies the “Area” in which a device is located, the actual host “ID,” and the destination application on that host, in the form of the “SEL” field. The CNLP address is logically segmented even further, as demonstrated by the following table:

IDP		DSP		
<i>AFI</i>	<i>IDI</i>	<i>HO-DSP</i>	<i>System-ID</i>	<i>NSEL</i>
Area Field		ID Field	SEL Field	

Observe the top row of the above figure. The ISO CLNP address provides granular control by separating *internal* and *external* routing information:

- The **IDP (Initial Domain Part)** portion of the address identifies the Autonomous System of the device (and is used to route *to* or *between* Autonomous Systems)
- The **DSP (Domain Specific Part)** portion of the address is used to route *within* the autonomous system.

* * *

All original material copyright © 2007 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

IS-IS Protocols and Addressing (continued)

IDP		DSP		
<i>AFI</i>	<i>IDI</i>	<i>HO-DSP</i>	<i>System-ID</i>	<i>NSEL</i>
Area Field		ID Field	SEL Field	

The IDP portion of the address is separated into two “sections,” including:

- **AFI (Authority and Format Identifier)** – specifies the organization authorized to assign addresses, and the format and length of the rest of the CLNP address. The AFI is always **8 bits**.
- **IDI (Initial Domain Identifier)** – identifies the “sub-organization” under the parent AFI organization. The length of the IDI is dependent on the chosen AFI.

An AFI of 0x49 indicates a *private* CLNP address, which cannot be routed globally (the equivalent of an IPv4 private address). An AFI of 0x47 is commonly used for global IS-IS networks, with the IDI section identifying specific organizations.

The AFI plus the IDI essentially identify the *autonomous system* of the address. However, this is not the equivalent of a BGP AS number, nor is it compatible with BGP as an exterior routing protocol.

The DSP portion of the address is separated into three “sections,” including:

- **HO-DSP (High Order DSP)** – identifies the *area* within an autonomous system
- **System ID** – identifies the specific host. Usually **48 bits** (or **6 octets**) in length, to accommodate MAC addresses
- **NSEL** – identifies the destination upper layer protocol of the host (always **8 bits**)

Two “types” of CLNP addresses are defined:

- **NET address** – does not contain upper-layer information (in other words, the SEL field is always set to 0x00)
- **NSAP address** – the “full” CLNP address, with populated Area, ID, and SEL fields.

Please note: A NET address is simply an NSAP address with a zero value in the SEL field.

* * *

All original material copyright © 2007 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

CLNS Address Example

The following is an example of a full ISO CLNS address:

47.1234.5678.9abc.def0.0001.1111.2222.3333.00

Correlating the above address to the appropriate fields:

<i>IDP</i>		<i>DSP</i>		
<i>AFI</i>	<i>IDI</i>	<i>HO-DSP</i>	<i>System-ID</i>	<i>NSEL</i>
47.	1234.5678.9abc.def0.	0001.	1111.2222.3333.	00
<i>Area Field</i>		<i>ID Field</i>		<i>SEL Field</i>

The *System-ID* is usually populated by the device's MAC address or IP v4 address.

Recall that CLNS addresses are of variable length. We can specify addresses without an *IDI* field:

47.0001.1111.2222.3333.00

Thus, the above address contains an *AFI* (Autonomous System), *HO-DSP* (Area), *System-ID* (in this example, a MAC Address), and the *NSEL* (SEL). Because the SEL field has a zero value (0x00), the above address is defined as a NET address, and not an NSAP address.

ISO CLNS addresses are not applied on an interface-by-interface basis. Instead, a single CLNS address is applied to the entire *device*.

Even if Integrated IS-IS is being used (thus indicating that IPv4 is being routed instead of CLNS), a CLNS address *is still required* on the IS-IS router. This is configured under the IS-IS router process.

Routers within the same area must share identical *AFI*, *IDI*, and *HO-DSP* values, but each must have a unique System-ID

All original material copyright © 2007 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

IS-IS Packet Types

IS-IS defines two categories of network devices:

- **ES (End System)** – identifies an *end* host.
- **IS (Intermediate System)** – identifies a Layer 3 router.

IS-IS additionally defines four categories of packet types:

- **Hello**
- **LSP**
- **CSNP**
- **PSNP**

Hello packets are exchanged for neighbor discovery. Three types of IS-IS Hello packets exist:

- **IIH (IS-IS Hello)** – exchanged between routers (or IS's) to form neighbor adjacencies.
- **ESH (ES Hello)** – sent from an ES to discover a router.
- **ISH (IS Hello)** – sent from an IS to announce its presence to ES's

An **LSP (Link State Packet)** is used to share topology information between routers. There are separate LSPs for Level 1 and Level 2 routing. LSP's are covered in great detail shortly.

A **CSNP (Complete Sequence Number PDU)** is an update containing the *full* link-state database. IS-IS routers will refresh the full database every **15 minutes**.

A **PSNP (Partial Sequence Number PDU)** is used by IS-IS routers to both *request* and *acknowledge* a link-state update.

* * *

All original material copyright © 2007 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

IS-IS Neighbors

IS-IS routers form neighbor relationships, called **adjacencies**, by exchanging **Hello** packets (often referred to as IS-IS Hellos or **IIH**'s). Hello packets are sent out every **10 seconds**, regardless of media type. Only after an adjacency is formed can routers share routing information.

IS-IS supports three IIH packet formats; one for **point-to-point** links, and two for **broadcast** (or **LAN**) links (Level-1 and Level-2 broadcast Hellos).

Unlike OSPF, IS-IS neighbors *do not* have to share a common IP subnet to form an adjacency. Adjacencies are formed across *CLNP* connections, not *IP* connections, even when using Integrated IS-IS. Thus, IS-IS actually requires *no* IP connectivity between its routers to route IP traffic!

There are two types of adjacencies:

- **Level-1 adjacency** – for routing *within* an area (intra-area routing)
- **Level-2 adjacency** – for routing *between* areas (inter-area routing)

IS-IS routers *must* share a common physical link to become neighbors, and the *System-ID* must be **unique** on each router. Additionally, the following parameters must be identical on each router:

- Hello packet format (point-to-point or broadcast)
- Hello timers
- Router “level” (explained shortly)
- Area (only for Level-1 adjacencies)
- Authentication parameters (Cisco devices currently support only **clear-text authentication** for IS-IS).
- MTU

Neighbors will elect a **DIS (Designated Intermediate System)** on broadcast links. A DIS is the equivalent of an OSPF DR (Designated Router). Unlike OSPF, however, there is no Backup DIS, and thus a new election will occur immediately if the DIS fails. Additionally, the DIS election is *preemptive*.

Whichever IS-IS router has the **highest priority** will be elected the DIS (default priority is **64**). In the event of a tie, whichever IS-IS router has the **highest SNPA** (usually **MAC**) **address** will become the DIS. The DIS sends out hello packets every **3.3 seconds**, instead of every 10 seconds.


(Reference: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcprt2/1cfisis.pdf)

* * *

All original material copyright © 2007 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

The IS-IS Hierarchy



IS-IS defines three types of IS-IS routers:

- **Level-1 Router** – contained within a single area, with a topology table limited to only its local area (called the **Level-1 Database**)
- **Level-2 Router** - a *backbone* router that routes *between* areas, and builds a **Level-2 Database**.
- **Level-1-2 Router** – similar to an *area border router*. Interfaces between a local area and the *backbone* area, and builds both a **Level-1** and a **Level-2** database.

Each type of IS-IS router will form only specific adjacencies:


- Level-1 routers form **Level-1 adjacencies** with other Level-1 routers and Level-1-2 routers.
- Level-2 routers form **Level-2 adjacencies** with other Level-2 routers and Level-1-2 routers.
- Level-1-2 routers form *both* **Level-1** and **Level-2 adjacencies** with other Level-1-2 routers.
- Level-1 routers ***will never*** form adjacencies with Level-2 routers.

The IS-IS backbone consists of multiple contiguous Level-2 routers, each of which can exist in a separate area.

* * *

All original material copyright © 2007 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

The IS-IS Hierarchy (continued)

Neighbors build their topology tables by sharing **LSP's (Link-State Packets)**, which are roughly the equivalent of OSPF LSA's. Depending on the type of adjacency, a router will send out either a Level-1 or Level-2 LSP.

Level-1 routers share Level-1 LSP's, and will build a Level-1 topology table consisting of *solely* its own area (thus forming the equivalent of an OSPF Totally Stubby area). If a Level-1 router has a packet destined for the *local* area, it simply routes the packet to the System ID by using the local topology table (Level-1 database).

If a Level-1 router has a packet destined for a *remote* area, it forwards it to the nearest Level-1-2 router. Level-1-2 routers set an **Attach (ATT) bit** in their Level-1 LSP's, informing other Level-1 routers that they are attached to another area.

Level-2 routers share Level-2 LSP's, and will build a Level-2 topology table, which contains a list of reachable areas across the IS-IS domain.

Level-1-2 routers will share *both* Level-1 *and* Level-2 LSP's with its appropriate adjacencies. Level-1-2 routers maintain separate Level-1 and Level-2 topology tables.

Level-1 routes (locally originated) are *always preferred* over Level-2 routes (originated from another area).

IS-IS routers will refresh the Link-State topology table every **15 minutes** (as opposed to every 30 minutes for OSPF).

(Reference: http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm)

* * *

All original material copyright © 2007 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

Basic IS-IS Configuration

To configure IS-IS, the IS-IS process must first be established:

```
Router(config)# router isis
```

The router must then be configured with a CLNP address:

```
Router(config)# router isis
Router(config-router)# net 49.0001.1921.6800.5005.00
```

To globally dictate the router-type of all interfaces (default is **level-1-2**):

```
Router(config)# router isis
Router(config-router)# is-type level-1
Router(config-router)# is-type level-1-2
Router(config-router)# is-type level-2
```

Finally, IS-IS must be explicitly enabled on the interface:

```
Router(config)# interface fa0/0
Router(config-if)# ip router isis
```

This not only allows IS-IS to form neighbor relationships out of this interface, it also adds the interface's network to the routing table.

The globally configured router-type can be overridden on each individual interface:

```
Router(config)# interface fa0/0
Router(config-if)# isis circuit-type level-1
Router(config-if)# isis circuit-type level-1-2
Router(config-if)# isis circuit-type level-2
```

To adjust the priority (default is **64**) of interface, increasing the likelihood that the router will be elected the DIS:


```
Router(config)# interface e0/0
Router(config-if)# isis priority 100
```

* * *

All original material copyright © 2007 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

IS-IS Passive-Interfaces



It is possible to control which router interfaces will participate in the IS-IS process. Just as with EIGRP and OSPF, we can use the *passive-interface* command.

However, please note that the *passive-interface* command works differently with IS-IS than with RIP or IGRP. IS-IS will no longer form neighbor relationships out of a “passive” interface, thus this command prevents updates from being *sent* or *received* out of this interface:

```

RouterC(config)# router isis
RouterC(config-router)# passive-interface s0
  
```

Router C will not form a neighbor adjacency with Router B.

We can configure **all** interfaces to be passive using the *passive-interface default* command, and then individually use the *no passive-interface* command on the interfaces we **do** want neighbors to be formed on:

```

RouterC(config)# router isis
RouterC(config-router)# passive-interface default
RouterC(config-router)# no passive-interface e0
  
```

Always remember, that the *passive-interface* command will prevent IS-IS (and OSPF) from forming neighbor relationships out of that interface. No routing updates are passed in either direction.

However, unlike OSPF, using the *passive-interface* command will still inject that interface’s network into the routing table. Thus, the *passive-interface* command can be useful when creating “stub” networks.

All original material copyright © 2007 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

The IS-IS Metric

IS-IS utilizes an arbitrary **cost** for its metric (the optional metrics of **delay**, **expense**, and **error** are not supported by Cisco). By default, interfaces of all types (regardless of speed) are assigned a metric of **10**.

To adjust the metric on an interface:

```
Router(config)# interface e0/0
Router(config-if)# isis metric 30
```

IS-IS Authentication

IS-IS authentication can be applied to a **link**, to an **area**, or to a **domain**. Remember, Cisco supports only clear-text authentication for IS-IS.

To configuration authentication on an interface-by-interface basis:

```
Router(config)# interface fa0/0
Router(config-if)# isis password MYPASSWORD level-1
Router(config-if)# isis password MYPASSWORD2 level-2
```

Note that separate authentication passwords can be applied to Level-1 or Level-2 Adjacencies. To configure authentication for an entire IS-IS area:

```
Router(config)# router isis
Router(config-router)# area-password MYPASSWORD
```

IS-IS Summarization

IS-IS supports both **inter-area** and **external** summarization, and uses the same command to accomplish both. If we wished to summarize the following networks into one summary route:

- 172.16.0.0/16
- 172.17.0.0/16
- 172.18.0.0/16
- 172.19.0.0/16
- 172.20.0.0/16
- 172.21.0.0/16
- 172.22.0.0/16
- 172.23.0.0/16

The following command would be required:

```
RouterC(config)# router isis
RouterC(config-router)# summary-address 172.16.0.0 255.248.0.0
```

* * *

All original material copyright © 2007 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

IS-IS and WAN Technologies

When configuring IS-IS over Frame-Relay, additional map statements are required:

```
Router(config)# interface s0/0
Router(config-if)# frame-relay map clns 105 broadcast
Router(config-if)# frame-relay map clns 106 broadcast
```

Additionally, we can map CLNP addresses in ISDN:

```
Router(config)# interface bri0
Router(config-if)# dialer map clns 49.0001.1921.6800.5005.00 name
MYNAME broadcast 3331111
```

IS-IS Troubleshooting

To view any CLNS neighbors, including the type of adjacency:

```
Router# show clns neighbors
```

To view only IS neighbors:

```
Router# show clns is-neighbors
```

To view specific IS-IS information about an interface:

```
Router# show clns interface e0/0
```

To view the IS-IS link-state topology table:

```
Router# show isis database
```

To view a list of all known IS-IS routers in all areas:

```
Router# show isis topology
```

(Reference: http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_command_reference_chapter09186a008017d02e.html)

* * *

All original material copyright © 2007 by Aaron Balchunas (aaron@routeralley.com),
unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

IS-IS vs. OSPF

IS-IS is often compared and contrasted to OSPF. Both protocols share several similarities, including:

- Both are *Link-State* routing protocols.
- Both use the *Dijkstra* algorithm to determine the shortest path.
- Both are *classless* and support *VLSMs*.
- Both use a *cost* metric.
- Both use *areas* to minimize the size of topology and routing tables.
- Both elect a *designated router* on broadcast links to contain link-state update traffic.

Despite these similarities, there are a multitude of crucial differences between IS-IS and OSPF, including:

- OSPF supports *only* IP, IS-IS supports *both* IP and CLNS.
- IS-IS *does not* require IP connectivity between routers to share routing information. Updates are sent via CLNS instead of IP.
- In OSPF, *interfaces* belong to areas. In IS-IS, the *entire router* belongs to an area.
- An IS-IS router belongs to only one Level-2 area, which results in less LSP traffic. IS-IS is thus more efficient and scalable than OSPF, and supports more routers per area.
- There is no Area 0 backbone area for IS-IS. The IS-IS backbone is a contiguous group of Level 1-2 and Level 2 routers.
- IS-IS *does not* elect a backup DIS. Additionally, DIS election is preemptive.
- On broadcast networks, even with an elected DIS, IS-IS routers still form adjacencies with *all* other routers. In OSPF, routers will only form adjacencies with the DR and BDR on broadcast links.
- IS-IS uses an arbitrary cost metric. OSPF's cost metric is based on the bandwidth of the link.
- IS-IS provides far more granular control of link-state and SPF timers than OSPF.

(Reference: <http://geocities.com/mnvhbatia/draft-bhatia-manral-diff-isis-ospf-00.txt>,
<http://www.ciscopress.com/articles/article.asp?p=31319&rl=1>)

* * *

All original material copyright © 2007 by Aaron Balchunas (aaron@routeralley.com),
unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.