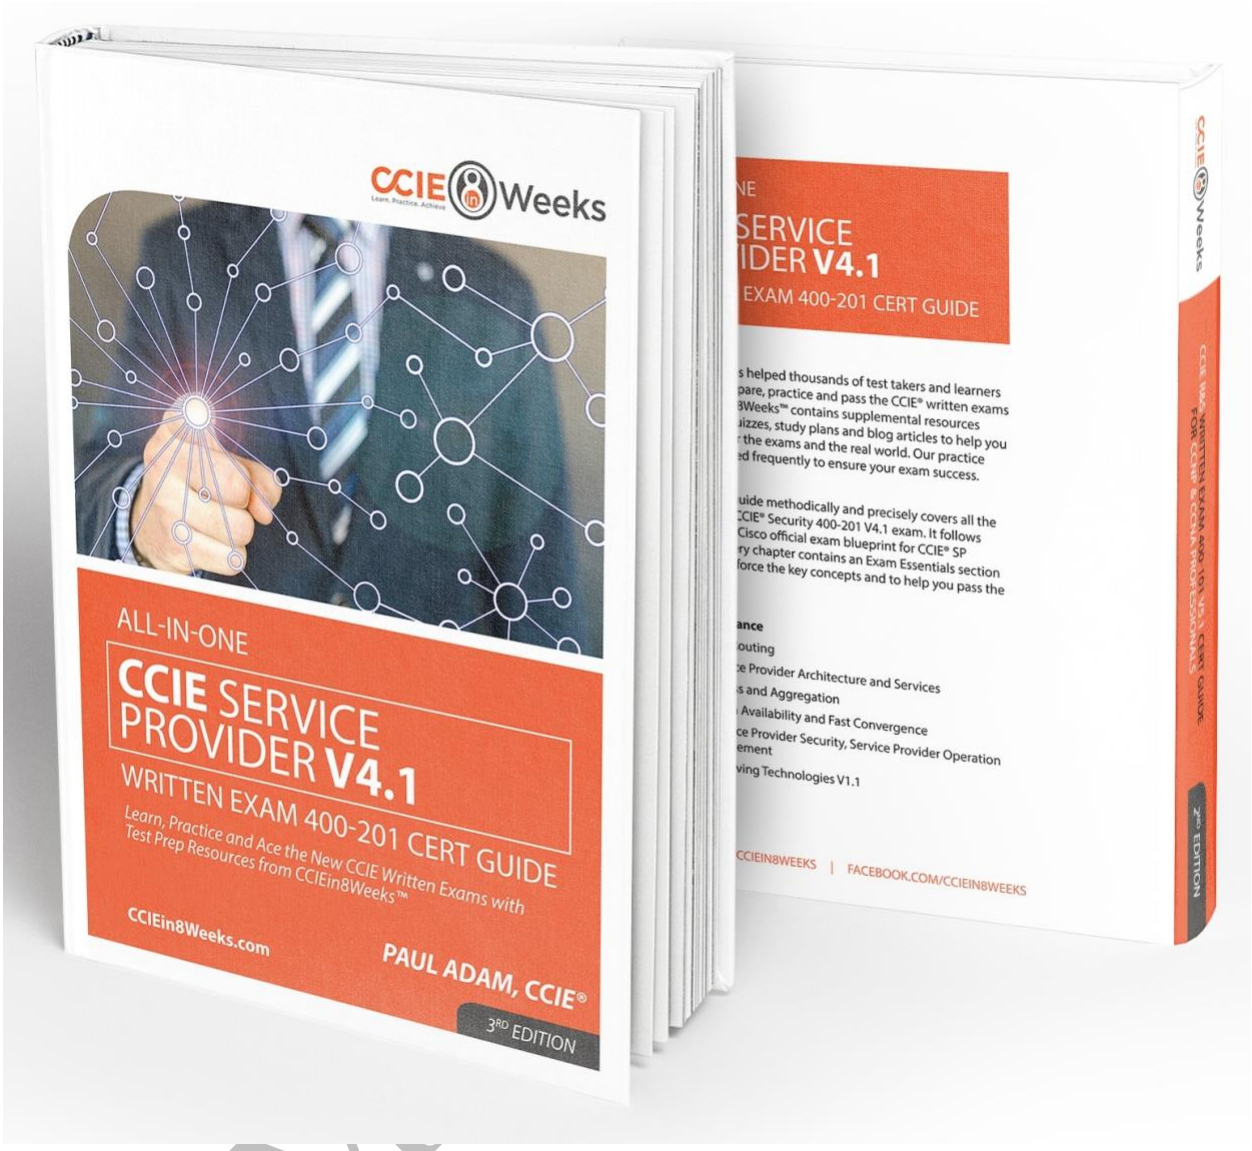
CCIE 8 Weeks
Learn. Practice. Achieve
ALL-IN-ONE
CCIE SERVICE PROVIDER V4.1
WRITTEN EXAM 400-201 CERT GUIDE
Learn, Practice and Ace the New CCIE Written Exams with Test Prep Resources from CCIEin8Weeks™
CCIEin8Weeks.com
PAUL ADAM, CCIE®
3RD EDITION
SERVICE
IDER V4.1
EXAM 400-201 CERT GUIDE
s helped thousands of test takers and learners
pare, practice and pass the CCIE® written exams
8Weeks™ contains supplemental resources
izzes, study plans and blog articles to help you
the exams and the real world. Our practice
ed frequently to ensure your exam success.
uide methodically and precisely covers all the
CIE® Security 400-201 V4.1 exam. It follows
Cisco official exam blueprint for CCIE® SP
ry chapter contains an Exam Essentials section
force the key concepts and to help you pass the
ance
outing
e Provider Architecture and Services
s and Aggregation
Availability and Fast Convergence
e Provider Security, Service Provider Operation
ement
ving Technologies V1.1
CCIEIN8WEEKS | FACEBOOK.COM/CCIEIN8WEEKS
CCIE 8 Weeks
CCIE WRITTEN EXAM 400-201 V3.1 CERT GUIDE FOR CCNP & CCNA PROFESSIONALS
2ND EDITION

All-in-One CCIE Service Provider V4.1

400-201 Written Exam Cert Guide

3rd Edition

Contents at a Glance

Part 6 Evolving Technologies V1.1

Table of Contents

Preface

Congratulations! You have taken your first step towards passing the CCIE Service Provider 400-201 (V4.1) written exam. This written exam cert guide along with practice quizzes from CCIEin8Weeks will help you prepare for this exam.

This book is dedicated to *all those souls who will never settle for less than they can be, do, share, and give!*

What this Exam Cert Guide covers

As you may already have noticed on the "Contents at a Glance" page that this guide has been formatted around the Cisco's official CCIE 400-201 exam topics or curriculum. So, as you read through parts and chapters, you know exactly where you're within your learning journey.
All contents are carefully compiled and covered in enough details however still trimmed to exactly what's necessary to pass the exam. The result of this approach is that you have 300 odd pages on your hands (as opposed to 700 to 1,400 pages long reference manuals!). Each exam topic has "Further Reading" section, which you can refer to if you are looking for more in-depth details or want to refer to the source of the original text. However, we have done the research for you in the form of curated set of links.

How to use this Exam Cert Guide

This guide is for anyone who's studying for CCIE Service Provider 400-201 (v4.1) written exam, regardless if this is your first time sitting for a CCIE written exam or you're a pro who's recertifying. I strongly suggest to take a methodical approach for exam preparation, i.e. start with a target date when you would like to sit for the actual exam and then work backwards to see what kind of study plan would work for you. I believe you can cover the entire contents within eight weeks including the practice questions (refer to website). We strongly suggest you to also use other study resources in addition to bringing your networking experience to bear in order to successfully pass the exam.

If you're totally in a crunch and consider yourself an advanced learner, you could try your luck by skimming through the "Exam Essentials" sections for each chapter- I guarantee you that you will be better prepared to tackle the exam with it than without!

What's available on the CCIEin8Weeks website

CCIEin8Weeks.com carries the extras (sold separately) that go hand in hand with this study guide to further ensure your exam success!

Website includes:

- Seven practice quizzes (one for each section as per official curriculum), One final exam
- Four study plans, to help you track your progress or help you come up with your own plan
- CCIE Exam community forum, to help you interact with others, share knowledge, find study partners etc.
- Last but not least, this exam cert guide in printable PDF format

Part 1 Core Routing

Chapter 1: Interior Gateway Protocol (IGP)

This chapter covers the following exam topics from Cisco's official 400-201 (v4.0) written exam curriculum.

- Describe, implement, and troubleshoot IS-IS
- Describe, implement, and troubleshoot OSPFv2 and OSPFv3
- Describe and optimize IGP scale and performance

Chapter 1: Interior Gateway Protocol (IGP)

Describe, implement, and troubleshoot IS-IS

Integrated Intermediate System-to-Intermediate System (IS-IS), Internet Protocol Version 4 (IPv4), is a standards-based Interior Gateway Protocol (IGP). Cisco IOS XR software implements the IP routing capabilities described in International Organization for Standardization (ISO)/International Engineering Consortium (IEC) 10589 and RFC 1995, and adds the standard extensions for single topology and multi topology IS-IS for IP Version 6 (IPv6).

IS-IS requires configuration on both the router and the interface. An IS-IS process is created when you enable IS-IS on a router and define a specific tag to identify that routing process. Interfaces configured with a specific tag will be part of the corresponding router process. More than one IS-IS process can run on a router for Connectionless Network Service (CLNS), but only one IS-IS process can run for IP. Small IS-IS networks are built as a single area that includes all the routers in the network. As the network grows larger, it is usually reorganized into a backbone area made up of the connected set of all Level 2 routers from all areas. The areas are connected to local areas. Within a local area, routers know how to reach all system IDs. Between areas, routers know how to reach the backbone, and the backbone routers know how to reach other areas.

Routers establish Level 1 adjacencies to perform routing within a local area (intra-area routing). Routers establish Level 2 adjacencies to perform routing between Level 1 areas (inter-area routing). If the network administrator does not specify Level 1 or Level 2 routing for the routing process being configured, the default routing behavior for the routing process will be Level 1-2. If Level 2 routing is configured on any process, additional processes are automatically configured as Level 1, with the exception of previously configured Level 2 process, which will remain Level 2. You can have only one Level-2 process. You can configure the Level-2 process to perform Level-1 routing at the same time. If Level-2 routing is not desired for a router instance, use the is-type command in router configuration mode to remove the Level-2 capability. You can also use the is-type command to configure a different router instance as a Level-2 router. Some networks use legacy equipment that supports only Level 1 routing. These

devices are typically organized into many small areas that cannot be aggregated due to performance limitations. Cisco routers are used to interconnect each area to the Level 2 backbone.
The idea behind the Designated Intermediate System (DIS) is similar to that behind the designated router in OSPF. The DIS creates a pseudo node (a virtual node), and all the routers on a LAN, including the DIS, form an adjacency with the pseudo node instead of forming n*(n-1) order adjacencies with each other in a full mesh.

On a LAN, one of the routers will elect itself the DIS based on interface priority (the default is 64). If all interface priorities are the same, the router with the highest subnetwork point of attachment (SNPA) is selected. MAC addresses are the SNPA on LANs. On Frame Relay networks, the local data-link connection identifier (DLCI) is the SNPA. If the SNPA is a DLCI and is the same at both sides of a link, the router with the higher system ID (in the NSAP address) will become the DIS.

A pseudo node LSP represents a LAN, including all ISs attached to that LAN, just as a non-pseudo node LSP represents a router, including all ISs and LANs connected with the router. The DIS election is pre-emptive (unlike with OSPF). If a new router boots on the LAN with a higher interface priority, it becomes the DIS, purges the old pseudo node LSP, and a new set of LSPs will be flooded. The DIS sends CSNPs describing all the LSPs in the database every 3 seconds. If a router needs an LSP because it is older than the LSP advertised by the DIS in its CSNP or it is missing an LSP that is listed in the CSNP, it will send a PSNP to the DIS and receive the LSP in return. This mechanism can work both ways: If a router sees that it has a newer version of an LSP, or it has an LSP that the DIS does not advertise in its CSNP, the router will send the newer or missing LSP to the DIS. The Cisco IS-IS implementation offers an authentication mechanism to prevent unauthorized routers from forming adjacencies or injecting TLVs. Currently, only plain-text authentication is available where the configured password is transmitted inside the IS-IS PDUs unencrypted in plain text. As such, the password can be determined by sniffing the packets. Future Cisco IOS Software releases will also contain Hashed Message Authentication Codes with MD5 (HMAC-MD5) with encrypted passwords as specified in the corresponding IETF draft.

IS-IS authentication is configured independently for adjacency establishment (hello) and for LSP authentication.

Single area, Single topology

Small IS-IS networks are built as a single area that includes all the routers in the network. As the network grows larger, it is usually reorganized into a backbone area made up of the connected set of all Level 2 routers from all areas, which is in turn connected to local areas. Within a local area, routers know how to reach all system IDs. Between areas, routers know how to reach the backbone, and the backbone routers know how to reach other areas.

Routers establish Level 1 adjacencies to perform routing within a local area (intra-area routing). Routers establish Level 2 adjacencies to perform routing between Level 1 areas (inter-area routing). Some networks use legacy equipment that supports only Level 1 routing. These devices are typically organized into many small areas that cannot be aggregated due to performance limitations. Cisco routers are used to interconnect each area to the Level 2 backbone.

Step 1	Define areas, prepare an addressing plan for the routers (including defining the NETs), and determine interfaces that will run Integrated IS-IS.
Step 2	Enable IS-IS as an IP routing protocol on the routers, and assign a tag to the process (if required).
Step 3	Configure the NETs on the routers. This identifies the routers for IS-IS.
Step 4	Enable Integrated IS-IS on the proper interfaces on the routers. Do not forget interfaces to stub IP networks, such as loopback interfaces (although there will not be any CLNS neighbors on these interfaces).

ISO has developed standards for two types of routing protocols:

- ES-IS discovery protocol—ES-IS performs "routing" between End Systems and Intermediate Systems referred as Level 0 "routing." ES-IS is analogous to the Address Resolution Protocol (ARP) in IP. Although it is not explicitly a routing protocol, ES-IS is included here because it is commonly used with routing protocols to provide end-to-end data movement through an internetwork.
- IS-IS routing protocols—IS-IS performs hierarchical (Level 1, Level 2, and Level 3) routing between intermediate systems. Level 3 routing is done between separate domains. However, note that the IS-IS routing protocol is not itself capable of Level 3 routing.

There is no Address Resolution Protocol (ARP), Internet Control Message Protocol (ICMP) or Inter-domain Routing Protocol (IDRP) for CLNS, but End System-to-Intermediate System (ES-IS) Protocol provides the same kind of reporting functions for ISs and ESs.

Further Reading

http://goo.gl/OqIRQP

Describe network types, levels and router types

NSAP addressing

NSAP is the network-layer address for CLNS packets. An NSAP describes an attachment to a particular service at the network layer of a node, similar to the combination of IP destination address and IP protocol number in an IP packet. NSAP encoding and format are specified by ISO 8348/Ad2.

ISO 8348/Ad2 uses the concept of hierarchical addressing domains. The global domain is the highest level. This global domain is subdivided into sub-domains, and each sub-domain is associated with an addressing authority that has a unique plan for constructing NSAP addresses. An NSAP address has two major parts: the initial domain part (IDP) and the domain specific part (DSP). The IDP consists of a 1-byte authority and format identifier (AFI) and a variable-length initial domain identifier (IDI), and the DSP is a string of digits identifying a particular transport implementation of a specified AFI authority. Everything to the left of the system ID can be thought of as the area address of a network node.

The LSP identifier is derived from the system ID (along with the pseudonode ID and LSP number). Each IS is usually configured with one NET and in one area; each system ID within an area must be unique.

Further Reading

http://goo.gl/s31fz4

Point-to-point, broadcast

When a network consists of only two networking devices connected to broadcast media and uses the integrated IS-IS protocol, it is better for the system to handle the link as a point-to-point link instead of as a broadcast link. This feature introduces a new command to make IS-IS behave as a point-to-point link between the networking devices.

Router(config-if)# isis network point-to-point | broadcast

Describe operations

From a high level, IS-IS operates as follows:

- Routers running IS-IS will send hello packets out all IS-IS-enabled interfaces to discover neighbors and establish adjacencies.
- Routers sharing a common data link will become IS-IS neighbors if their hello packets contain information that meets the criteria for forming an adjacency. The criteria differ slightly depending on the type of media being used (p2p or broadcast). The main criteria are matching authentication, IS-type and MTU size.
- Routers may build a link-state packet (LSP) based upon their local interfaces that are configured for IS-IS and prefixes learned from other adjacent routers.
- Generally, routers flood LSPs to all adjacent neighbors except the neighbor from which they received the same LSP. However, there are different forms of flooding and also a number of scenarios in which the flooding operation may differ.
- All routers will construct their link-state database from these LSPs.
- A shortest-path tree (SPT) is calculated by each IS, and from this SPT the routing table is built.

Configuration Steps

On Cisco equipment there are some basic configuration parameters which are required to get IS-IS up and running. The commands to get IS-IS up and running on a router include:

1. router isis – This command is used to enable the start of the IS-IS process on the router
2. net address – This command is used to configure the NET address which will be used to identify this router
3. ip router isis – This command is configured on each interface which will participate with IS-IS

By default, a Cisco router will operate in Level-1/Level-2 mode which supports both levels of routing. This makes IS-IS easier to configure but requires additional memory and processor on each device. If only Level-1 or Level-2 routing is required on a specific device you can enable only the level required on the device by using the is-type [level-1 | level-2] command while in IS-IS router configuration mode.

The sample configurations below configure all the routers in the above topology with these parameters:

- Area 49.0001
- Level 1 (L1) and Level 2 (L2) routers (this is the default unless otherwise specified)
- No optional parameters
- Running IS-IS for IP only
- Loopback interfaces (loopbacks are advertised by IS-IS, not IS-IS enabled)

Router 1

```
!
interface Loopback0
ip address 172.16.1.1 255.255.255.255

!--- Creates loopback interface and assigns
!--- IP address to interface Loopback0.

!
interface Ethernet0
ip address 172.16.12.1 255.255.255.0
ip router isis


 !--- Assigns IP address to interface Ethernet0
 !--- and enables IS-IS for IP on the interface.

!
router isis
passive-interface Loopback0
net 49.0001.1720.1600.1001.00
!

!--- Enables the IS-IS process on the router,
 !--- makes loopback interface passive
 !--- (does not send IS-IS packets on interface),
 !--- and assigns area and system ID to router.
```

Router 2

```
!
interface Loopback0
ip address 172.16.2.2 255.255.255.255

!--- Creates loopback interface and assigns
 !--- IP address to interface Loopback0.

!
Interface Ethernet0
ip address 172.16.12.2 255.255.255.0
ip router isis

!--- Assigns IP address to interface Ethernet0
 !--- and enables IS-IS for IP on the interface.

!
Interface Serial0
ip address 172.16.23.1 255.255.255.252
ip router isis

!--- Assigns IP address to interface Serial0
 !--- and enables IS-IS for IP on the interface.

!
router isis
 passive-interface Loopback0
net 49.0001.1720.1600.2002.00
!

!--- Enables the IS-IS process on the router,
!--- makes loopback interface passive
 !--- (does not send IS-IS packets on interface),
 !--- and assigns area and system ID to router.
```

Router 3
! interface Loopback0 ip address 172.16.3.3 255.255.255.255 !--- Creates loopback interface !--- and assigns IP address to !--- interface Loopback0. ! Interface Serial0 ip address 172.16.23.2 255.255.255.252 ip router Isis !--- Assigns IP address to !--- interface Serial0 and enables !--- IS-IS for IP on the interface. ! router isis passive-interface Loopback0 net 49.0001.1234.1600.2231.00 ! !--- Enables the IS-IS process on the router, !--- makes loopback interface passive !--- (does not send IS-IS packets on interface), !--- and assigns area and system ID to router.

Describe optimization features

To minimize the number of adjacencies, LSDBs, and related SPF and PRC computations that are performed, it is recommended that you have configured all Level 1 routers as Level 1 by using the is-type command. We recommend that you use the metric-style wide command because some

features, such as setting prefix tags and MPLS traffic engineering, require that routers that are running IS-IS generate the new-style TLVs that have wider metric fields. If you use the default narrow metric style for IS-IS, the router generates and accepts old-style type, length, and value objects (TLVs).

Metrics, wide metric

Cisco IOS Software allows for wide metrics with the support of a 24-bit metric field. Using the new metric style, link metrics now have a maximum value of 16777215 (2^{24}-1) with a total path metric of 4261412864 (254 x 2^{24}). Deploying IS-IS in the IP network with wide metrics is recommended to enable finer granularity and to support future applications such as Traffic Engineering.

Running different metric styles within one network poses a serious problem: Link-state protocols calculate loop-free routes because all routers (within one area) calculate their routing table based on the same link-state database. This principle is violated if some routers look at old-style (narrow), and some at new-style (wider) TLVs. However, if the same interface cost is used for both the old- and new-style metrics, then the SPF will compute a loop-free topology.

Monitoring IS-IS Adjacencies

Use the show clns neighbor command to display the adjacencies for a specific router. This is the output of this command from Router 1 (R1) and Router 2 (R2):

```
R1# show clns neighbor
System Id   Interface  SNPA        State  Holdtime  Type Protocol
R2          Et0        0000.0c47.b947  Up    24       L1L2 ISIS

R2# show clns neighbor
System Id  Interface  SNPA           State  Holdtime  Type Protocol
R1         Et0        0000.0c09.9fea   Up    24       L1L2 ISIS
R3         Se0        *HDLC*           Up    28       L1L2 ISIS
```

In the above example, R1 recognizes R2 on its E0 interface with the adjacency type being L1L2. Because R1 and R2 are configured with default configurations, they send and receive both L1 and L2 hellos.

R2 recognizes R1 on its E0 interface, and Router 3 (R3) on its S0 interface. The same explanation as above is true for the adjacency type. Since R1 and R2 are on the same Ethernet

interface, there is a DIS for both L1 and L2. You can verify this using the show clns interface <int> command on Router 1, as shown below:

```
R1# show clns interface ethernet 0
Ethernet0 is up, line protocol is up
  Checksums enabled, MTU 1497, Encapsulation SAP
  Routing Protocol: ISIS
    Circuit Type: level-1-2
    Interface number 0x0, local circuit ID 0x1
    Level-1 Metric: 10, Priority: 64, Circuit ID: R2.01
    Number of active level-1 adjacencies: 1
    Level-2 Metric: 10, Priority: 64, Circuit ID: R2.01
    Number of active level-2 adjacencies: 1
    Next ISIS LAN Level-1 Hello in 5 seconds
    Next ISIS LAN Level-2 Hello in 1 seconds
```

In the above output, R2 is the DIS. It is the R2 (DIS) that generates the pseudonode Link State Packet (LSP) and is denoted with a non-zero LSP-ID - R2.01 Since the Metric/Priority are the same for both routers in L1/L2, the tiebreaker for the DIS is the highest Subnetwork Points of Attachment (SNPA) address on the LAN segment. The SNPA address refers to the data link address, and in this case is the MAC address.

Notice that the DIS is elected for both levels, and that no backup DIS exists, as with Open Shortest Path First (OSPF), which has a backup Designated Router (DR). Some other points of interest from the above output include:

- Circuit Type: L1L2
- L1 and L2 metrics and priorities are at default values: 10 and 64
- L1 and L2 adjacencies: 1 (from R1 perspective on the Ethernet interface - it is R2 only)
- IS-IS LAN hellos for L1 and L2
- Maximum Transmission Unit (MTU): 1497. This is because the Open Systems Interconnection (OSI) IS-IS header is encapsulated inside a 3 byte 802.2 header.

Monitoring the IS-IS Database

The show isis database (detail) command displays the contents of the IS-IS database. This is the output of this command when issued on R2. Since IS-IS is a link state protocol, the link state database should be the same for any router in the same area.

```
R2# show isis database
ISIS Level-1 Link State Database:
LSPID         LSP Seq Num  LSP Checksum  LSP Holdtime    ATT/P/OL
R1.00-00      0x0000008B   0x6843        55              0/0/0
R2.00-00    * 0x00000083   0x276E        77              0/0/0
R2.01-00    * 0x00000004   0x34E1        57              0/0/0
R3.00-00      0x00000086   0xF30E        84              0/0/0
ISIS Level-2 Link State Database:
LSPID         LSP Seq Num  LSP Checksum  LSP Holdtime    ATT/P/OL
R1.00-00      0x00000092   0x34B2        41              0/0/0
R2.00-00    * 0x0000008A   0x7A59        115             0/0/0
R2.01-00    * 0x00000004   0xC3DA        50              0/0/0
R3.00-00      0x0000008F   0x0766        112             0/0/0
```

There are a few things to notice in the above output. First, about the LSP-ID:
The LSP-ID, R1.00-00, can be broken down into three sections: R1/00/00

- R1 = system ID
- 00 = non-zero value for the pseudonode. Notice R2.01-00 is the pseudonode LSP.
- 00 = fragment number. In this case, there are only fragment numbers of 00, which indicates that all the data fit into this LSP fragment, and there was no need to create more fragments. If there had been information that did not fit into the first LSP, IS-IS would have created more LSP fragments, such as 01, 02, and so on.

The * denotes the LSPs that were generated by this router, the router that the show command was issued on. Also, since this router is an L1 and L2 router, it contains an L1 and L2 database. You can also look at a specific LSP and use the detail keyword to show more information. An example of this is shown here:

```
R2# show isis database R2.00-00 detail
ISIS Level-1 LSP R2.00-00
LSPID         LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
```

```
R2.00-00    * 0x00000093   0x077E       71          0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  Hostname: R2
  IP Address:   172.16.2.2
  Metric: 10        IP 172.16.12.0 255.255.255.0
  Metric: 0         IP 172.16.2.2 255.255.255.255
  Metric: 10        IP 172.16.23.0 255.255.255.252
  Metric: 10        IS R2.01
  Metric: 10        IS R3.00
ISIS Level-2 LSP R2.00-00
LSPID         LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
R2.00-00    * 0x0000009A   0x5A69        103           0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  Hostname: R2
  IP Address:   172.16.2.2
  Metric: 10        IS R2.01
  Metric: 10        IS R3.00
  Metric: 10        IP 172.16.23.0 255.255.255.252
  Metric: 10        IP 172.16.1.1 255.255.255.255
  Metric: 10        IP 172.16.3.3 255.255.255.255
  Metric: 0         IP 172.16.2.2 255.255.255.255
  Metric: 10        IP 172.16.12.0 255.255.255.0
```

The above output shows that the loopback address of this router is advertised with a value of 0. This is because the loopback is advertised with a passive-interface command under the router IS-IS process, and the loopback interface by itself is not enabled for IS-IS. All other IP prefixes have a value of 10, which is the default cost on the interfaces running IS-IS.

Further Reading

http://goo.gl/wlJP2g

http://goo.gl/aPyHGG

http://goo.gl/iRYLmC

http://goo.gl/W623T4

Describe, implement, and troubleshoot OSPFv2 and OSPFv3

LSA types (1, 2, 3, 4, 5, 7, 9)

Table 1-1, shows LSA types and their meaning

LSA Type	Description
1	Router Link advertisements. Generated by each router for each area it belongs to. They describe the states of the router's link to the area. These are only flooded within a particular area.
2	Network Link advertisements. Generated by Designated Routers. They describe the set of routers attached to a particular network. Flooded in the area that contains the network.
3 or 4	Summary Link advertisements. Generated by Area Border routers. They describe inter-area (between areas) routes. Type 3 describes routes to networks, also used for aggregating routes. Type 4 describes routes to ASBR.
5	AS external link advertisements. Originated by ASBR. They describe routes to destinations external to the AS. Flooded all over except stub areas.
7	Routers in a Not-so-stubby-area (NSSA) do not receive external LSAs from Area Border Routers, but are allowed to send external routing information for redistribution. They use type 7 LSAs to tell the ABRs about these external routes, which the Area Border Router then translates to type 5 external LSAs and floods as normal to the rest of the OSPF network.

9	a link-local "opaque" LSA (defined by RFC2370) in OSPFv2 and the Intra-Area-Prefix LSA in OSPFv3. It is the OSPFv3 LSA that contains prefixes for stub and transit networks in the link-state ID.

A designated router (DR) is the router interface elected among all routers on a particular multi-access network segment, generally assumed to be broadcast multi-access. The basic neighbor discovery process (Hello), flooding (224.0.0.6), DR election (priority, RID). Special techniques, often vendor-dependent, may be needed to support the DR function on non-broadcast multi-access (NBMA) media. It is usually wise to configure the individual virtual circuits of a NBMA subnet as individual point-to-point lines. DRs exist for the purpose of reducing network traffic by providing a source for routing updates. The DR maintains a complete topology table of the network and sends the updates to the other routers via multicast. All routers in a multi-access network segment will form a slave/master relationship with the DR. They will form adjacencies with the DR and BDR only. Every time a router sends an update, it sends it to the DR and BDR on the multicast address 224.0.0.6. The DR will then send the update out to all other routers in the area, to the multicast address 224.0.0.5. This way all the routers do not have to constantly update each other, and can rather get all their updates from a single source. The use of multicasting further reduces the network load. DRs and BDRs are always setup/elected on OSPF broadcast networks. DR's can also be elected on NBMA (Non-Broadcast Multi-Access) networks such as Frame Relay. DRs or BDRs are not elected on point-to-point links (such as a point-to-point WAN connection) because the two routers on either sides of the link must become fully adjacent and the bandwidth between them cannot be further optimized. DR and non-DR routers evolve from 2-way to full adjacency relationships by exchanging DD, Request, and Update.

Table 1-2, shows network types and recommended traffic type for use

Network Type	Traffic Type
non-broadcast p2mp broadcast	Unicast
broadcast p2p	Multicast

p2mp	
loopback	Stub

Route types (N1, N2, E1, E2)

There are several types of OSPF routes:

- Intra-Area—In a multi-area OSPF network, routes, originated within an area, are known by the routers in the same area as Intra-Area routes. These routes are flagged as O in the show ip route command output.
- Inter-Area—When a route crosses an OSPF Area Border Router (ABR), the route is known as an OSPF Inter-Area route. These routes are flagged as O IA in the show ip route command output.
- Both Intra and Inter-Area routes are also called OSPF Internal routes, as they are generated by OSPF itself, when an interface is covered with the OSPF network command.
- External Type-2 or External Type-1—Routes which were redistributed into OSPF, such as Connected, Static, or other Routing Protocol, are known as External Type-2 or External Type-1. These routes are flagged as O E2 or O E1 in the show ip route command output.
- NSSA external type 2 or NSSA external type 1—When an area is configured as a Not-So-Stub Area (NSSA), and routes are redistributed into OSPF, the routes are known as NSSA external type 2 or NSSA external type 1. These routes are flagged as O N2 or O N1 in the show ip route command output.

Implement and troubleshoot neighbor relationship

You can use the show ip ospf neighbor command to determine the state of the OSPF neighbor or neighbors.

If the show ip ospf neighbor command reveals nothing at all—or reveals nothing about the particular neighbor you are analyzing—then this router has not seen any “valid” OSPF HELLOs from that neighbor. This means that OSPF either did not receive any HELLO packets from the neighbor or received HELLO packets that failed very basic sanity checks.

- state = down

A neighbor that is discovered dynamically through reception of HELLO packets can fall back to a down state if it is being deleted, for example when OSPF does not receive HELLO packets from the neighbor for period of time longer than the Dead timer interval. Therefore, the down

state is transient for such neighbors; they will either advance to higher states or be completely deleted from the table of known neighbors. This is known as being "forgotten".
Usually, neighbors that are seen in the down state were manually configured with the neighbor command. Manually configured neighbors are always present in the OSPF neighbor table. If OSPF has never received HELLO packet from the manually configured neighbor, or if no HELLO packets were heard from the neighbor during the previous Dead timer interval, then the manually configured neighbor will be listed as down.

- state = init
  The init state indicates that a router sees HELLO packets from the neighbor, but two-way communication has not been established. A Cisco router includes the Router IDs of all neighbors in the init (or higher) state in the Neighbor field of its HELLO packets. For two-way communication to be established with a neighbor, a router also must see its own Router ID in the Neighbor field of the neighbor's HELLO packets.
- state = exstart / exchange
  OSPF neighbors that are in exstart or exchange state are trying to exchange DBD packets. The router and its neighbor form a master and slave relationship. The adjacency should continue past this state. If it does not, there is a problem with the DBD exchange, such as a maximum transmission unit (MTU) mismatch or the receipt of an unexpected DBD sequence number
- state = 2-way
  The 2-way state indicates that the router has seen its own Router ID in the Neighbor field of the neighbor's HELLO packet. Receiving a Database Descriptor (DBD) packet from a neighbor in the init state will also a cause a transition to 2-way state. The OSPF neighbor 2-way state is not a cause for concern (e.g. for broadcast network type).
- state = loading
  In the loading state, routers send link-state request packets. During the adjacency, if a router receives an outdated or missing link-state advertisement (LSA), it requests that LSA by sending a link-state request packet. Neighbors that do not transition beyond this state are most likely exchanging corrupted LSAs. This problem is usually accompanied by a %OSPF-4-BADLSA console message.

A common problem when using Open Shortest Path First (OSPF) is routes in the database don't appear in the routing table. In most cases OSPF finds a discrepancy in the database so it doesn't install the route in the routing table. Often, you can see the Adv Router is not-reachable message

(which means that the router advertising the LSA is not reachable through OSPF) on top of the link-state advertisement (LSA) in the database when this problem occurs. Here is an example:

Router# show ip ospf database router 172.16.32.2
Adv Router is not-reachable
LS age: 418
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 172.16.32.2
Advertising Router: 172.16.32.2
LS Seq Number: 80000002
Checksum: 0xFA63
Length: 60
Number of Links: 3

There are several reasons for this problem, most of which deal with mis-configuration or a broken topology. When the configuration is corrected the OSPF database discrepancy goes away and the routes appear in the routing table.

Further Reading

http://goo.gl/WmL2EB
http://goo.gl/A8N2zE

Implement and troubleshoot OSPFv3 address-family support

Configuration Steps

1. enable
2. configure terminal
3. router ospfv3 [process-id]
4. area area-ID [default-cost | nssa | stub]
5. auto-cost reference-bandwidth Mbps
6. bfd all-interfaces
7. default {area area-ID[range ipv6-prefix | virtual-link router-id]} [default-information originate [always | metric | metric-type | route-map] | distance | distribute-list prefix-list prefix-list-name {in | out} [interface] | maximum-paths paths | redistribute protocol | summary-prefix ipv6-prefix]

8. ignore lsa mospf
9. interface-id snmp-if-index
10. log-adjacency-changes [detail]
11. passive-interface [default | interface-type interface-number]
12. queue-depth {hello | update} {queue-size | unlimited}
13. router-id {router-id}

To filter outgoing link-state advertisements (LSAs) to an Open Shortest Path First (OSPF) interface, use the ip ospf database-filter all out command in interface or virtual network interface configuration modes. To restore the forwarding of LSAs to the interface, use the no form of this command.
Router(config-if)#ipv6 ospf neighbor 2002:234::2
OSPFv3:Neighbor address needs to be a link-local address

You need to add the additional static frame-relay DLCI mapping statement to neighbor's link local address before OSPF adjacency will form.
ip ospf database-filter all out [disable]

Verification Steps

- show ospfv3 [process-id] border-routers
- show ospfv3 [process-id [area-id]] database [database-summary | internal | external[ipv6-prefix] [link-state-id] | grace | inter-area prefix [ipv6-prefix | link-state-id] | inter-area router [destination-router-id | link-state-id] | link [interface interface-name |link-state-id] | network [link-state-id] | nssa-external [ipv6-prefix] [link-state-id] | prefix [ref-lsa {router | network} | link-state-id] | promiscuous | router [link-state-id] | unknown [{area | as | link} [link-state-id]] [adv-router router-id] [self-originate]
- show ospfv3 [process-id] events [generic | interface | lsa | neighbor | reverse | rib | spf]
- show ospfv3 [process-id] [area-id] flood-list interface-type interface-number
- show ospfv3 [process-id] graceful-restart
- show ospfv3 [process-id] [area-id] interface[type number] [brief]
- show ospfv3 [process-id] [area-id] neighbor[interface type interface-number] [neighbor-id] [detail]
- show ospfv3 [process-id] [area-id] request-list[neighbor] [interface] [interface neighbor]
- show ospfv3 [process-id] [area-id] retransmission-list [neighbor] [interface] [interface neighbor]

- show ospfv3 [process-id] statistic[detail]
- show ospfv3 [process-id] summary-prefix
- show ospfv3 [process-id] timers rate-limit
- show ospfv3 [process-id] traffic[interface-type interface-number]
- show ospfv3 [process-id] virtual-links

Further Reading

http://goo.gl/N8Z90m

IPv4/v6 address-family

The OSPFv3 address families feature enables both IPv4 and IPv6 unicast traffic to be supported. With this feature, users may have two router processes per interface, but only one process per AF. If the IPv4 AF is used, an IPv4 address must first be configured on the interface, but IPv6 must be enabled on the interface. A single IPv4 or IPv6 OSPFv3 process running multiple instances on the same interface is not supported.

Users with an IPv6 network that uses OSPFv3 as its IGP may want to use the same IGP to help carry and install IPv4 routes. All routers on this network have an IPv6 forwarding stack. Some (or all) of the links on this network may be allowed to do IPv4 forwarding and be configured with IPv4 addresses. Pockets of IPv4-only routers exist around the edges running an IPv4 static or dynamic routing protocol. In this scenario, users need the ability to forward IPv4 traffic between these pockets without tunneling overhead, which means that any IPv4 transit router has both IPv4 and IPv6 forwarding stacks (e.g., is dual stack). This feature allows a separate (possibly incongruent) topology to be constructed for the IPv4 AF. It installs IPv4 routes in IPv4 RIB, and then the forwarding occurs natively. The OSPFv3 process fully supports an IPv4 AF topology and can redistribute routes from and into any other IPv4 routing protocol.

An OSPFv3 process can be configured to be either IPv4 or IPv6. The address-family command is used to determine which AF will run in the OSPFv3 process, and only one address family can be configured per instance however multiple instances of OSPFv3 can be enabled on single interface. Once the AF is selected, users can enable multiple instances on a link and enable address-family-specific commands.

Different instance ID ranges are used for each AF (it is carried inside packet header as instance ID field and that's unique to IPv6). Each AF establishes different adjacencies, has a different link state database, and computes a different shortest path tree. The AF then installs the routes in AF-

specific RIB. LSAs that carry IPv6 unicast prefixes are used without any modification in different instances to carry each AFs' prefixes.
The IPv4 subnets configured on OSPFv3-enabled interfaces are advertised through intra-area prefix LSAs, just as any IPv6 prefixes. External LSAs are used to advertise IPv4 routes redistributed from any IPv4 routing protocol, including connected and static. The IPv4 OSPFv3 process runs the SPF calculations and finds the shortest path to those IPv4 destinations. These computed routes are then inserted in the IPv4 RIB (computed routes are inserted into an IPv6 RIB for an IPv6 AF).

Because the IPv4 OSPFv3 process allocates a unique pdbindex in the IPv4 RIB, all other IPv4 routing protocols can redistribute routes from it. The parse chain for all protocols is same, so the ospfv3 keyword added to the list of IPv4 routing protocols causes OSPFv3 to appear in the redistribute command from any IPv4 routing protocol. With the ospfv3 keyword, IPv4 OSPFv3 routes can be redistributed into any other IPv4 routing protocol as defined in the redistribute ospfv3 command. In case of using OSPF for IP VPNs as PE-CE protocol, the down bit helps prevent routing loops between MP-BGP and OSPF, but not when external routes are announced, such as when redistribution between multiple OSPF domains or when external routes are injected in an area that is dual-homed to the provider network. The PE router redistributes an OSPF route from a different OSPF domain into an OSPF domain as an external route. The down bit is not set because LSA Type 5 does not support the down bit. The redistributed route is propagated across the OSPF domain.