



CCIE Routing and Switching Written Exam Study Guide

for the CCIE Routing and Switching Written Exam version 3.0



Email

sales@ccbootcamp.com

Phone

1.877.NLI.CCIE (654.2243)

Int'l: +1 702.968.5100

Website

www.ccbootcamp.com

Forums

www.routerie.com

www.securityie.com

www.voiceie.com



CCBOOTCAMP'S CCIE Routing and Switching Written Exam Study Guide

for the CCIE Routing and Switching Written Exam version 3.0

For questions about this workbook please visit: www.routerie.com

CCBOOTCAMP

375 N. Stephanie Street
Building 21, Suite 2111
Henderson, NV 89014
1.877.654.2243 Toll Free

www.ccbootcamp.com

"Cisco," the "Cisco Logo," "CCNA," "CCNP," "CCDP," "CCDA," "CCIE," "Cisco Certified Network Associate," "Cisco Certified Design Professional," "Cisco Certified Design Associate," "and "Cisco Certified Network Professional," are registered trademarks of Cisco Systems, Inc. The contents contained wherein, is not associated or endorsed by Cisco Systems, Inc.

PLEASE READ THIS SUBSCRIPTION LICENSE AGREEMENT CAREFULLY BEFORE USING THIS PRODUCT. THIS SUBSCRIPTION LICENSE AGREEMENT APPLIES TO **CCBOOTCAMP's CCIE Routing and Switching Written Exam Study Guide**.

BY ORDERING THIS PRODUCT YOU ARE CONSENTING TO BE BOUND BY THIS LICENSING AGREEMENT. IF YOU DO NOT AGREE TO ALL OF THE TERMS OF THIS LICENSE, THEN DO NOT PURCHASE THIS PRODUCT.

License Agreement

CCBOOTCAMP's CCIE Routing and Switching Written Exam Study Guide is copyrighted. In addition, this product is at all times the property of CCBOOTCAMP, and the customer shall agree to use this product only for themselves, the licensed user. The license for the specific customer remains valid from the purchase date until they pass their CCIE Routing and Switching written exam.

CCBOOTCAMP's CCIE Routing and Switching Written Exam Study Guide is licensed by individual customer. This material cannot be resold, transferred, traded, sold, or have the price shared in any way. Each specific individual customer must have a license to use this product. The customer agrees that this product is always the property of CCBOOTCAMP, and they are just purchasing a license to use it. A Customer's license will be revoked if they violate this licensing agreement in any way.

Copies of this material in any form or fashion are strictly prohibited. If for any reason a licensed copy of this material is lost or damaged a new copy will be provided free of charge, except for the cost of printing, shipping and handling.

Individuals or entities that knowingly violate the terms of this licensing agreement may be subject to punitive damages that CCBOOTCAMP could seek in civil court. Damages will be limited to a maximum of \$500,000.00 per individual and \$2,000,000.00 per entity. In addition, individuals or entities that knowingly violate the terms of this license agreement may be subject to criminal penalties as are allowed by law.

The venue of any dispute, controversy, litigation or proceeding (formal or informal) arising out of or pertaining to this licensing agreement or the subject hereof shall lie exclusively in the County of Clark, State of Nevada. Provided, however, that if any such dispute, controversy, litigation or proceeding requires or permits jurisdiction in a federal court or agency of the United States, then venue shall lie in no federal court or agency other than those located in (or nearest to) the County of Clark, State of Nevada.

Term and Termination of License Agreement

This License is effective until terminated. Customer may terminate this License at any time by destroying all copies of written and electronic material of said product. Customer's rights under this License will terminate immediately without notice from CCBOOTCAMP, if Customer fails to comply with any provision of this License. Upon termination, Customer must destroy all copies of material in its possession or control. The license for the specific user remains valid from the purchase date until the user passes their lab exam pertaining to the purchased subscription. Once the customer passes the relevant lab exam the license is terminated and all material written or electronic in their possession or control must be destroyed or returned to CCBOOTCAMP.

Warranty

No warranty of any kind is provided with this product. There are no guarantees that the use of this product will help a customer pass any exams, tests, or certifications, or enhance their knowledge in any way. The product is provided on an "AS IS" basis. In no event will CCBOOTCAMP, its suppliers, or licensed resellers be liable for any incurred costs, lost revenue, lost profit, lost data, or any other damages regardless of the theory of liability arising out of use or inability to use this product.

About the author:

Author – Brad Ellis

Brad Ellis, co-founder and the CEO of Network Learning, Inc., a Cisco value-added reseller, and its renowned training subsidiary, CCBOOTCAMP, began with a simple vision – to take the training and professional development of Cisco Certified networking engineers to a whole new level. Brad also wanted to create a one-stop shop that could sell, install and service Cisco networking gear with a superior network engineering staff, while offering individuals and organizations the opportunity to obtain Cisco Certification for their existing engineers and IT staff. Over the past decade, Brad, along with his team of CCBOOTCAMP instructors, have helped more than 2000 network engineers reach their goal of becoming CCIE certified, and countless businesses obtain their networks, service and professional training, all from one trusted source. Brad currently maintains the following designations: CCIE #5796 (Routing and Switching, Security), CCSI #30482, CSS1, CCDP, CCNP, MCNE, MCSE.

Table of Contents

Disclaimer	i
Trademark Acknowledgements	i
Feedback Information	ii
Chapter 1 General Networking Theory.....	1
General Routing Concepts	1
Distance-Vector Routing Protocols	1
Comparison of routing algorithms	2
Link-State Routing Protocols	2
Hybrid Routing Protocols	3
Routing Protocol Concepts	4
Routing Loops.....	5
Route Summarization	5
Classfull and a Classless routing protocol	8
Routing Information Base (RIB) and Routing Protocols Interaction.....	8
Routing Decision Criteria	8
Administrative Distance.....	9
Routing Table	9
RIB and Forwarding Information Base interaction.....	10
Redistribution between routing	10
Route-maps	12
Prefix List.....	14

Table of Contents (Continued)

Access Control Lists (ACLs).....	15
Distribution Lists	15
Filter Lists.....	16
Routing Loops and Split Horizon	17
Troubleshooting routing loop	18
Chapter 1 Questions.....	19
Chapter 1 Answers	27
Chapter 2 Bridging & LAN Switching	28
Trunking.....	29
ISL vs. 802.1Q	30
Catalyst 6500 Switched Port Analyzer (SPAN) Information	31
Gigabit Ethernet.....	32
Dynamic Inter-Switch Link Protocol (DISL)	32
Fast EtherChannel (FEC)	32
Using PAgP to Configure Fast EtherChannel.....	32
Link Aggregate Control Protocol (LACP)	34
Virtual LAN (VLAN)	34
VLAN Trunk Protocol (VTP)	35
VTP message types:	36
Spanning-Tree.....	37
Spanning-Tree Protocol (STP) 802.1d.....	37
Types of STP on Cisco Switches.....	37
Redundancy Without Loops.....	38
Root Bridges and Switches	39

Table of Contents (Continued)

Bridge Protocol Data Units (BPDUs)	39
Operation Under STP	41
STP Step-by-Step.....	42
STP Timers.....	43
Port State Progression in STP.....	43
802.1w Rapid Spanning Tree Port States	43
Summary Table	44
802.1w Rapid Spanning Tree Port Roles.....	44
802.1w vs. 802.1d	44
STP Broadcast Domain Characteristics.....	44
Topology Changes—TCN, TCA, and TC.....	44
STP Enhancements.....	45
Unidirectional Link Detection (UDLD).....	46
Loop Guard versus UDLD.....	46
Spanning Tree PortFast BPDUs Guard	47
Traffic Suppression (Storm Control)	47
Unicast Flooding.....	47
Cisco Spanning Tree Root Guard.....	48
Multicast	49
Ip Multicasting MAC address mapping	49
Ethernet Issues.....	50
Duplex mismatch error.....	50
Chapter 2: Questions.....	51
Chapter 2 Answers	57

Table of Contents (Continued)

Chapter 3 Internet Protocol (IP)	58
IP Addressing	58
Subnetting	59
Subnetting Tricks	59
Route Summarization	60
Ports and Sockets	61
Network Address Translation (NAT)	61
CIDR and VLSM	62
Hot Standby Router Protocol (HSRP)	63
Virtual Router Redundancy Protocol (VRRP)	64
Gateway Load Balancing Protocol (GLBP)	64
Network Time Protocol (NTP)	65
Introduction to NTP	65
NTP Design Criteria	67
NTP Architecture	70
Clock Technology and Public Time Servers	71
The ntp server Global Configuration Command	72
IP Services	72
BOOTP	74
BOOTP Procedure	74
Dynamic Host Configuration Protocol (DHCP)	74
Cisco IOS DHCP Server	75
IP Applications	75
WCCP	76

Table of Contents (Continued)

Network Management	76
SNMPv1 Operations	77
SNMPv2 additional Operations (also has SNMPv1 Ops).....	78
SNMPv3	78
SNMP Communities	78
SNMP Traps and Notifications.....	79
SNMP Embedded Event Manager	79
Syslogs.....	80
Chapter 3 Questions	81
Chapter 3 Answers	86
Chapter 4 IP Routing Protocols	87
Routing Information Protocol (RIP) & RIP V2	87
Split Horizon in a Hub and Spoke Network	87
Open Shortest Path First (OSPF).....	89
Other OSPF Features:	89
OSPF Traffic Types:	89
OSPF Metrics	89
Passive OSPF Interface	90
OSPF Multicast Addresses	90
Default Routes	90
OSPF Timers.....	90
Init	93
Enhanced Interior Gateway Routing Protocol (EIGRP)	96
Types of EIGRP Successors	98

Table of Contents (Continued)

Feasibility Condition	98
Attributes of EIGRP	98
EIGRP Tables.....	98
Choosing routes.....	99
Init Flag.....	100
EIGRP Stub Routing.....	101
Simple Hub and Spoke Network	102
Route Summary.....	103
Auto-Summarization.....	103
Process ID for an Autonomous System	103
Show IP Route EIGRP	103
Show Ip Eigrp Topology	104
Show Ip Eigrp Neighbor	106
Border Gateway Protocol (BGP).....	107
Situations that may require BGP:	107
Interior Border Gateway Protocol (IBGP)	107
Exterior Border Gateway Protocol (EBGP)	108
BGP Attributes	108
Weight Attribute	108
Local Preference Attribute	109
Multi-Exit Discriminator Attribute.....	110
Origin Attribute.....	111
AS_path Attribute	112
Next-Hop Attribute	113
Community Attribute	113

Table of Contents (Continued)

Cluster-List	113
Originator ID	114
BGP Neighbor Connectivity	114
Synchronization/Full Mesh	115
Next-Hop-Self Command	115
Private AS numbers	115
BGP Path Selection	116
Scalability Problems with Internal BGP (IBGP)	116
Peer Groups	116
Confederations	117
Route Reflectors.....	117
Route Summary	117
BGP Clusters	118
Route Maps	118
No Export.....	118
Route Dampening.....	118
Backdoor	119
Enabling BGP Routing	119
Controlling BGP Routes	119
Policy-Based Routing	120
Policy-Based Routing Benefits	120
Data Forwarding Using Policy-Based Routing	121
Tagging Network Traffic	121
Applying Policy-Based Routing.....	121
Policy Route Maps	122

Table of Contents (Continued)

Match Clauses Define the Criteria	122
Set Clauses Define the Route	123
Source-Sensitive and Equal-Access Routing	123
Quality of Service (QoS).....	124
Load Sharing	124
Management Implications.....	124
PBR Summary	125
Route Filtering	125
IP Prefix-list	126
OSPF ABR Type 3 LSA Filtering.....	127
Configuring OSPF ABR Type 3 LSA Filtering	127
The use of SHOW and DEBUG commands	128
Chapter 4 Questions	129
Chapter 4 Answers	140
Chapter 5 Quality of Service (QoS)	141
QoS Overview.....	142
Five Benefits for Implementing QoS in the Enterprise Networks	142
How a Converged Network Behaves Without QoS.....	142
QoS framework.....	142
Call Admission Control Functionality.....	143
Integrated Services vs. Differentiated Services.....	143
Configure QoS Policy using Modular QoS CLI	145
Classification and Marking	150
Purposes of Classification and Marking	150
Difference between Classification and Marking.....	150

Table of Contents (Continued)

Class of Service, IP Precedence and DiffServ Code Points	151
Network Based Application Recognition (NBAR)	152
Classify and Mark Traffic	153
Congestion Management	155
Identify and Differentiate Between IOS Queuing Techniques.....	155
Apply Each Queuing Technique to the Appropriate Application.....	157
IP RTP Priority and Low Latency Queuing (LLQ) Differences	158
Configure WFQ, CBWFQ, and LLQ	159
Congestion Avoidance.....	161
Explain How TCP Responds to Congestion	161
Explain Tail Drop and Global Synchronization	161
Identify and Differentiate Between: RED, WRED, FRED	162
Configure IOS Congestion Avoidance Features.....	162
Link Efficiency Tools	163
The Need for Link Efficiency Tools.....	163
Real Time Protocol Header Compression (CRTP)	165
Configure and Monitor Various LFI methods and CRTP	166
Policing and Shaping.....	169
The Difference between Policing and Shaping and How Each Relates to QoS	169
When to Apply and How to Configure Policing Mechanisms	169
Different Types of Traffic Shaping and How to Apply Them.....	170
Configure the Different Types of Traffic Shaping	171
First-In, First-Out (FIFO)	175

Table of Contents (Continued)

Weighted Fair Queuing (WFQ)	175
Priority Queuing	175
Custom Queuing	177
Class-Based Weighted Fair Queuing	181
Packet over SONET/SDH (PoS) and IP Precedence	181
IP Precedence	182
Random Early Detection (RED)	183
Weighted Random Early Detection (WRED)	183
Weighted Round-Robin (WRR)/Queue Scheduling	183
Class of Service (CoS)	185
Shaping vs. Policing	185
Traffic Shaping	185
Committed Access Rate (CAR)	187
Network-Based Application Recognition (NBAR)	188
Configuring NBAR	189
Differentiated Services Code Point (DSCP)	190
Resource Reservation Protocol (RSVP)	191
Load Balancing	192
802.1x and QoS	193
Syntax	193
Custom Queuing (CQ)	194
Why Use CQ?	194
Restrictions	195
Configuring a Traffic Policy	196
Attaching a Traffic Policy to an Interface	196

Table of Contents (Continued)

Configuring a Traffic Class with NBAR Example	197
ToS Byte	198
DiffServ Field	198
Differences between Traffic-Shaping Mechanisms.....	199
CQ and Extended Burst Capability.....	201
Committed Access Rate (CAR) definition	202
Analysis	204
Connecting from Spoke to Spoke.....	206
Chapter 5 Questions	207
Chapter 5 Answers	213
Chapter 6 Wide Area Networking (WAN)	214
Leased Line Protocols.....	214
High-Level Data Link Control (HDLC).....	214
Point-to-Point Protocol (PPP)	215
Modems and Async.....	215
Frame Relay.....	216
Basic Facts	216
Types of Circuits	216
Data Link Connection Identifier (DLCI)	216
Local Management Interface (LMI)	217
DLCI Capacity Calculation.....	217
Encapsulation	218
Split Horizon and Frame Relay Interfaces.....	218
Speed Elements	219
Congestion	219

Table of Contents (Continued)

Frame Relay Compression	220
Frame-Relay Mapping	220
Other Frame Relay Issues	221
Frame Relay Adaptive Traffic Shaping	221
Physical Layer	223
Serial Interface Abbreviations	223
Is Your Interface a DTE or a DCE?	223
RS-232/EIA-232	224
V.35 Interface	224
Troubleshooting Serial Links	226
Show Controllers Command	227
Debug Commands	228
Increasing Output Drops	229
Increasing Input Drops	230
Excessive Aborts	231
Proper steps to resolve abort problems:	231
Clocking Problems	231
Increasing Interface Resets on a Serial Link	232
Increasing Carrier Transitions Count on Serial Link	233
CRC and Framing Errors	233
To resolve CRC and Framing error problems:	233
Alarms	234
Receive Alarm Indication Signal (Blue)	234
Receive Remote Alarm Indication (Yellow)	234

Table of Contents (Continued)

Transmit Sending Remote Alarm (Red)	234
Transmit Remote Alarm Indication (Yellow)	234
Transmit Alarm Indication Signal (Blue).....	234
Dynamic Packet Transport/Spatial Reuse Protocol (DPT/SRP)	234
Data Compression	234
Stacker Compression	235
Predictor Compression	235
Dynamic Multipoint VPN (DMVPN)	235
Routing with DMVPN	236
Configuration requirements for DMVPN	236
Chapter 6 Questions	237
Chapter 6 Answers	241
Chapter 7 IP Multicast.....	242
Benefits of IP Multicast.....	242
Multicast	242
IGMP and CGMP Multicast Protocols	243
Designated Querier.....	244
IGMP Versions 1, 2, and 3	245
Multicast Forwarding and Distribution Trees	250
Rendezvous Points (Auto-RP, BSR)	250
Recommended Rendezvous Point Placement	251
Group-RP Mapping Mechanism	251
Comments on Auto-RP	252

Table of Contents (Continued)

Comments on Static RP.....	252
Calculating a Multicast Address	253
Protocol Independent Multicast (PIM)	253
PIM Commands.....	255
Reverse Path Forwarding (RPF)	255
PIM and Distance Vector Multicast Routing Protocol (DVMRP)	255
PIM-SM Mechanics (Joining, Pruning PIM State, Mroute table).....	256
PIM-DM	257
Bidirectional PIM (bidir-PIM)	258
Designated Forwarder (DF) Election	260
Bidirectional Group Tree Building.....	260
Packet Forwarding.....	261
Memory, Bandwidth, and CPU Requirements	261
Debugging bidir-PIM is easier than PIM-SM.....	262
RP Tree Delivery for All Packets.....	262
Bidir-PIM Partial Upgrades Not Allowed.....	262
Bidir-PIM Network Redundancy Not Supported	262
Bidir-PIM Nonbroadcast Multiaccess Mode Not Supported.....	262
Bidir-PIM Traffic Forwarding Restrictions.....	262
Chapter 7 Questions.....	268
Chapter 7 Answers	271
Chapter 8 Security	272
Bridge/Switch Security.....	272
Understanding How Port Security Works.....	272

Table of Contents (Continued)

Allowing Traffic Based on Host MAC Address	272
Restricting Traffic Based on Host MAC Address	273
Guidelines for Port Security Configuration	273
MAC Address Access-Lists (ACL)	274
Protocol Type-Code Access-Lists (ACL)	275
VLAN Access-list (VACL)	277
Private VLANs	278
802.1x	279
Access Lists	279
Sample Standard Access List	281
Sample Extended Access-list	281
Sample Named Access-list	282
IOS Firewall (CBAC)	283
RADIUS and TACACS+	284
TCP and UDP	284
Packet Encryption Differences	284
Authentication and Authorization	284
Multiprotocol Support	285
Router Management	285
Interoperability	285
Traffic	286
TACACS+ Traffic Example	286
RADIUS Traffic Example	287
AAA Security Services	288
AAA Philosophy	289

Table of Contents (Continued)

Benefits of Using AAA	290
Method Lists.....	290
The First Step, or Where to Begin	291
Overview of the AAA Configuration Process	291
Unicast RPF	291
SMURF Attack.....	292
IP Spoofing	293
Non-Blind Spoofing.....	293
Blind Spoofing	293
Man in the Middle Attack	294
Denial of Service Attack	295
Chapter 8 Questions	296
Chapter 8 Answers	299
Chapter 9 MPLS	300
Transport.....	300
IP/Multi Protocol Label Switching (IP/MPLS)	300
MPLS Overview	302
Forwarding Equivalence Class (FEC).....	302
Architectural Blocks of MPLS	303
MPLS Virtual Private Networks	309
VPN Operation	309
VPN Route Target Communities.....	310
MPLS Forwarding	311
Chapter 9: Questions.....	313

Table of Contents (Continued)

Chapter 9 Answers	315
Chapter 10 IPV6	316
Internet Protocol Version 6 (IPv6)	316
Unchanged characteristics of Addressing in IPv6	316
Addressing	316
Zero Compression in IPv6 Addresses.....	317
IPv6 Mixed Notation	317
IPv6 Address Prefix Length Representation	317
IPv6 Address Types	317
Important IPv6 address blocks	318
Aggregatable Global Addresses	318
Site-Local Addresses (Also known as Unique)	318
Link-Local Addresses	318
IPv6 Multicast Addresses	319
IPv6 Multicast Address Format	319
IPV6 Multicast.....	320
IPv6 Anycast Addresses	320
IPv6 neighbor discovery protocol.....	320
Host-Router Discovery Functions	321
Host-Host Communication Functions	321
Redirect Function	321
IPv6 ND Functions Compared to Equivalent IPv4 Functions.....	322
Host-Router Discovery Functions Performed By Routers	322
Host-Router Discovery Functions Performed By Hosts	323

Table of Contents (Continued)

Next-Hop Determination	323
Address Resolution	323
Duplicate Address Detection	324
IPv6 Tunneling	324
IPv6 Manually Configured Tunnels	325
IPv6 over IPv4 GRE Tunnels	325
Automatic 6to4 Tunnels	325
Automatic IPv4-Compatible IPv6 Tunnels.....	326
The Intra-Site Automatic Tunnel Addressing Protocol (ISATAP Tunnels)	326
OSPFv3 vs. OSPFv2	326
LSA Types for IPv6	327
NBMA in OSPF for IPv6	328
Importing Addresses into OSPF for IPv6	328
Multicast Listener Discovery Protocol for IPv6.....	332
Chapter 10 Questions	333
Chapter 10 Answers	335

Introduction - *YOUR JOURNEY IS JUST BEGINNING!*

By Brad Ellis

I got my first CCIE certification back in April of 2000. That was quite an exciting time! There were no lab boot camps, no training classes, and there was VERY little training material available to use for practice (except, of course, www.ccbootcamp.com). Since then, I've spent many hours answering e-mails, phone calls, and posts on message forums in an attempt to help other CCIE candidates find their shortest path to becoming a CCIE. The CCIE process is two fold, the written exam which qualifies you for your lab exam attempt, and the actual lab exam itself. Both of these steps need to be broken down into smaller steps and appropriately attacked. The first big step, the written exam, is all theory. Meaning you need to understand the meat and bones behind the Cisco networking technologies that are going to be presented to you in this book. Having a good understanding of these technologies will enable you to properly prepare for the second big (or shall we say, HUGE) step, the lab exam.

When I prepared for my written exam, I was forced to use many books and try to ascertain what was applicable to the exam, and what wasn't. I probably over studied by a couple weeks for the exam, but I'm sure it paid off when I prepared for my lab exam! My goal, in authoring this book, was to collect ALL of the relevant information and put it together in a one-stop-shop book. While this book is enough for you to pass the written exam, it is certainly not enough for you to use to pass the lab exam. There are many other books available for the lab exam which goes into much greater detail on the subjects covered in this book. We'll save the lab exam preparation for another time (and my next book!).

Why are there wireless questions on the written exam while wireless is not present in the lab exam? Could this be a hint of things to come? Only time will tell!

One of the biggest reasons why folks tend to start studying for the CCIE exams and never actually accomplish their goal is due to (1) time and (2) focus. Set your goals appropriately, try to finish reading this book within four to six weeks, and plan on taking your written exam after that. Then, schedule your lab exam for six to nine months out, and start studying for that. Put aside plenty of time to tackle the written and lab exams. Set the proper expectations with your family members. Let's get this first step out of the way together, and then you can focus on overcoming the lab exam!

Good luck on the first step of your networking journey!

Chapter 1

General Networking Theory

General Routing Concepts

Distance-Vector Routing Protocols

DV protocols periodically pass full copies of their routing tables to all of their immediate neighbors. Each recipient then increments the values, updates its routing table, and forwards that out through all of its interfaces to its neighboring routers. These routing protocols understand the direction and distance to any network connection on the internetwork. These updates are sent at specific increments, usually every 30 to 90 seconds, and contain the entire routing table.

Once this information has made the rounds, each router will have built a routing table with information about the "distances" to networked resources. It does not learn anything specific about other routers, or the network's actual topology. Because the "distance" usually depends on the number of "hops" to a destination, network distances and costs are rough estimates that do not relate directly to either physical distances or speed of the intermediate links.

The primary benefit of DV protocols is that they are easy to configure and maintain. For this reason they are quite common in small networks that have few redundant paths, and no stringent network performance requirements. The most common DV routing protocol is Routing Information Protocol (RIP), which uses a single distance metric (hops) to determine the best next path to take for any given packet. Cisco's proprietary Interior Gateway Routing Protocol (IGRP) would be another example of a DV routing protocol.

In any internetwork with redundant routes, using a dynamic routing protocol is better than using static routes, because the routing protocol will have the flexibility to automatically detect and correct failures in the network.

The problems associated with DV protocols include slow convergence, routing loops, counting to infinity, and excessive bandwidth utilization.

RIP version 1 and RIP version 2 are examples of distance-vector routing protocols.

A distance-vector routing protocol uses the Bellman-Ford algorithm to calculate paths. Examples of distance-vector routing protocols include RIPv1 or 2 and IGRP. EGP and BGP are not pure distance-vector routing protocols but their

concepts are the same. In many cases, EGP and BGP are considered DV (distance-vector) routing protocols. A distance-vector routing protocol requires that a router informs its neighbors of topology changes periodically and, in some cases, when a change is detected in the topology of a network. Compared to link-state protocols, which requires a router to inform all the nodes in a network of topology changes, distance-vector routing protocols have less computational complexity and message overhead.

Comparison of routing algorithms

Distance-vector routing protocols are simple and efficient in small networks, and require little, if any management. However, naïve distance-vector algorithms do not scale well (due to the count-to-infinity problem), and have poor convergence properties, which has led to the development of more complex but more scalable algorithms, such as link-state routing protocols and loop-free distance-vector protocols, for use in large networks.

The primary advantage of link-state routing is that it reacts more quickly, and in a bounded amount of time, to connectivity changes. Also, the link-state packets that are sent over the network are smaller than the packets used in distance-vector routing. Distance-vector routing requires a node's entire routing table to be transmitted, while in link-state routing only information about the node's immediate neighbors are transmitted. Therefore, these packets are small enough that they do not use network resources to any significant degree. The primary disadvantage of link-state routing is that it requires more storage and more computing to run than distance-vector routing.

Link-State Routing Protocols

Link-State Routing Protocols develop and maintain a full knowledge of the network's routers as well as how they connect to one another. This information is gathered through the exchange of *link-state advertisements* (LSAs) between routers. The LSAs are used to develop a topological database, which the shortest path algorithm then uses to compute reachability to networked destinations. This process allows quick discovery of changes in the network topology, either because of a component failure, or as a result of changes by the network engineer.

One of the biggest advantages to Link-State protocols is that they avoid the problem of wasted bandwidth that comes from DV routing protocols sending out their full routing tables several times a minute. On a properly configured network, this will leave more bandwidth available for passing user traffic.

Other advantages to Link-State routing protocols include:

- Faster convergence
- Greater scalability, allowing bigger, more robust networks
- Changes in topology can be sent out immediately, so convergence can be quicker

- They take bandwidth into account when determining routes

The concerns with Link-State protocols include:

- During the initial discovery process, link-state routing protocols can flood the network, decreasing the network's capability to transport data.
- Link-state routing is both memory and processor intensive.

OSPF and ISIS are examples of Link State protocols. The link-state protocol is performed by every switching node in the network (i.e. nodes which are prepared to forward packets; in the Internet, these are called routers). The basic concept of link-state routing is that every node receives a map showing connectivity of the network, in the form of a graph showing which nodes are connected to which other nodes.

Each node then independently calculates the best next hop from it for every possible destination in the network. (It does this using only its local copy of the map, and without communicating in any other way with any other node.) The collection of best next hops forms the routing table for the node. This contrasts with distance-vector routing protocols, which work by having each node share its routing table with its neighbors. In a link-state protocol, the only information passed between the nodes is information used to construct the connectivity maps.

For a link-state routing protocol:

- Each router transmits routing information to all nodes in the flooding domain
- Each router knows of all other routers in the flooding domain
- Each router builds an individual picture of the flooding domain
- Distance-vector routing protocols

Hybrid Routing Protocols

A more recent development is that of loop-free distance-vector protocols, e.g. EIGRP. Such protocols are as robust and manageable as distance-vector protocols, while avoiding counting to infinity and hence having good worst-case convergence times.

Hybrid Routing Protocols take into account basic distance-vector metrics, but also incorporate other more accurate metrics in their calculations. They converge more rapidly than distance-vector protocols, while avoiding the processing overhead associated with link-state updates. Also, they are event driven rather than using a timer to decide when to send updates; this conserves bandwidth for the transmission of user data.

Cisco's proprietary Enhanced Interior Gateway Routing Protocol (EIGRP) is the most common hybrid routing protocol. It was designed to combine the best

aspects of distance-vector and link-state routing protocols without incurring any of the performance limitations specific to either. Remember that one of the major limitations to EIGRP is that it only runs on Cisco equipment.

Routing Protocol Concepts

Convergence—The process of bringing the routing tables on all the routers in the network to a consistent state. Different routing protocols will converge at different rates depending on their design. Link State protocols will usually converge faster than Distance Vector protocols. *Convergence time* is how long it takes for all the routers in a given system to share information.

Load Balancing—Allows the transmission of packets to a specific destination over two or more paths. This can depend on equal or unequal cost; and can be configured per packet or per destination. Quite often it takes some thought and manual configuration to achieve the desired result.

Static Routing—The information in a router's route table can be built manually through static route entries, or dynamically through a routing protocol. Static routes can point to a specific host, a network, a subnet, or a super-net. You can also have floating static routes; routes that have an Administrative Distance (AD) set higher than the in-use dynamic routing protocol. If the route learned through the dynamic routing protocol is lost, then the floating static route will come into play. This provides a pre-configured automatic fallback route.

IOS command to add a static route - **ip route 192.168.10.0 255.255.255.0 192.168.1.1**

With a floating static route - **ip route 192.168.10.0 255.255.255.0 192.168.1.1 105 ***

**The last argument of "105" on the ip route command is the administrative distance assigned to the route.*

Metrics—All routing protocols use metrics to calculate the best path. Some protocols use simple metrics, such as RIP which uses hop count. Others, such as EIGRP, use more meaningful information. Other metrics that you may encounter include load, delay, reliability and cost. Sometimes system administrators will manually configure the metrics on a router to control the routing behavior of their network.

Route Flapping—The frequent changing of preferred routes as an interface or router goes into and out of operation (error condition). This process can create problems in a network, especially in complex OSPF networks, as this information will cause the routers to constantly recalculate their OSPF database and flood the network with LSAs.

Autonomous Systems (ASs)—A group of routers sharing a single routing policy, which run under a single technical administration, and commonly with a single Interior Gateway Protocol (IGP). Each AS has a unique identifying number

between 1 and 65,535 (64,512 through 65,535 are set aside for private use) usually assigned by an outside authority. Passing routing information between ASs is performed through an exterior gateway protocol, such as BGP.

Route Tagging—provides the capability to have flexible policy controls by creating a 32-bit tag value specific to the local routing domain to advertise to external routes. This is particularly useful in transit ASs, where the IGP interacts with BGP.

Periodic Updates—Routing protocols that use this technique send their routing updates at a consistent update interval, typically anywhere from 10 to 90 seconds. If the link speed is slow, or if the link is DDR, this process can present problems in some networks. If the update interval is too high, then the convergence speed of the network may suffer.

Passive-Interface—Prevents interfaces from sending routing updates. They will, however, continue to listen for updates. This command is applied in the router configuration, and specifies a physical interface.

Routing Loops

Routing loops occur when the routing tables of some or all of the routers in a given domain route a packet back and forth without ever reaching its final destination. Routing loops often occur during route redistribution, especially in networks with multiple redistribution points.

There are several commonly used methods for preventing routing loops, including:

- **Holddowns**—Routes are held for a specified period of time to prevent updates advertising networks that are possibly down. The period of time varies between routing protocols, and is configurable. Holddown timers should be set very carefully—if they are too short, they are ineffective; too long and convergence will be delayed.
- **Triggered updates**—Also known as flash updates, are sent immediately when a router detects that a metric has changed or a network is no longer available. This helps speed convergence. Instead of waiting for a certain time interval to elapse to update the routing tables, the new information is sent as soon as it is learned.
- **Split horizon**—If a router has received a route advertisement from another router, it will not re-advertise it back out the interface from which it was learned.
- **Poison reverse**—After once learning of a route through an interface, advertise it as unreachable back through that same interface.

Route Summarization

Route summarization condenses routing information by consolidating like routes, and collapsing multiple subnet routes into a single network route.

Where summarization is not applied, each router in a network must retain a route to every subnet in the network. This means as the network grows, the routing table becomes larger and larger. Routers that have had their routes summarized can reduce some sets of routes to a single advertisement, which reduces the load on the router and simplifies the network design.

For instance, let's consider a router that has several interfaces that have the following addresses:

- Interface s0 - 172.16.215.0/24
- Interface s1 - 172.16.126.0/24
- Interface s2 - 172.16.227.0/24
- Interface s3 - 172.16.218.0/24
- Interface s4 - 172.16.219.0/24
- Interface s5 - 172.16.129.0/24
- Interface s6 - 172.16.119.0/24
- Interface s7 - 172.16.117.0/24

Provided this address sequence was not used elsewhere on the network, an upstream neighbor could summarize these addresses as 172.16.0.0/16 and have only a single route in its table.

For another example, consider that you had a router with interfaces configured as follows:

- Interface s0 - 172.108.168.0/24
- Interface s1 - 172.108.169.0/24
- Interface s2 - 172.108.170.0/24
- Interface s3 - 172.108.171.0/24
- Interface s4 - 172.108.172.0/24
- Interface s5 - 172.108.173.0/24

The entire range of subnets could be summarized, as 172.108.168.0/21 and an upstream neighbor would only have to maintain a single route in its table.

Let's take one more example, but this time review the actual bits involved:

- a. 172.16.25.0/24
- b. 172.16.26.0/24
- c. 172.16.27.0/24
- d. 172.16.28.0/24
- e. 172.16.29.0/24
- f. 172.16.30.0/24

First let's translate the decimal values of the IP addresses to binary:

- a. 10101100.00010000.00011001.00000000
- b. 10101100.00010000.00011010.00000000
- c. 10101100.00010000.00011011.00000000
- d. 10101100.00010000.00011100.00000000
- e. 10101100.00010000.00011101.00000000
- f. 10101100.00010000.00011110.00000000

Now let's compare and determine the least significant digit where the numbers remain identical:

- | | | |
|----|-------------------------|--------------|
| a. | 10101100.00010000.00011 | 001.00000000 |
| b. | 10101100.00010000.00011 | 010.00000000 |
| c. | 10101100.00010000.00011 | 011.00000000 |
| d. | 10101100.00010000.00011 | 100.00000000 |
| e. | 10101100.00010000.00011 | 101.00000000 |
| f. | 10101100.00010000.00011 | 110.00000000 |

You have just discovered the summary address and subnet: 172.16.24.0/21

Some important reasons to take advantage of summarization:

- The larger the routing table, the more memory is required because every entry takes up some of the available memory.
- The routing decision process may take longer to complete as the number of entries in the table are increased.
- An added benefit of reducing the IP routing table size is that it requires less bandwidth and time to advertise the network to remote locations, thereby increasing network performance.

For large networks, the reduction in route propagation and routing information overhead can be significant. Route summarization is of minor concern in production networks until their size gets considerable. However, if summarization has not been taken into account during the initial design phase, it is very difficult to implement later.

Some routing protocols, such as EIGRP, summarize automatically. Other routing protocols, such as OSPF, require manual configuration to support route summarization.

A routing protocol can summarize on a bit boundary only if it supports *variable-length subnet masks* (VLSMs).

Remember that when redistributing routes from a routing protocol that supports VLSM (such as EIGRP or OSPF) into a routing protocol that does not (such as RIPv1 or IGRP) you might lose some routing information.

Most specific network match is used first for a router running multiple protocols to learn how to reach a destination network/host.