ALL-IN-ONE
Data Center Core Technologies (DCCOR) V1.0 Exam

CCIE and CCNP DC Core 350-601 Exam Cert Guide

1st Edition

Copyright © 2020 CCIEin8Weeks.com
All rights reserved.

Sample Guide

Our Cisco Next-Level Certification Catalog

FullStackNetworker
1st EDITION
FullStackNetworker.com
ALL-IN-ONE
DEVELOPING APPLICATIONS USING CISCO CORE PLATFORMS AND APIS
Cisco DevNet Professional DEVCOR 350-901 Core Exam
Learn, Practice and Ace the New Cisco DevNet Exams with Test Prep Resources from Full Stack Networker™
Muhammad Afaq Khan, CCIE #9070
ALL-IN-ONE DEVELOPING APPLICATIONS USING CISCO CORE PLATFORMS AND APIS
1ST EDITION

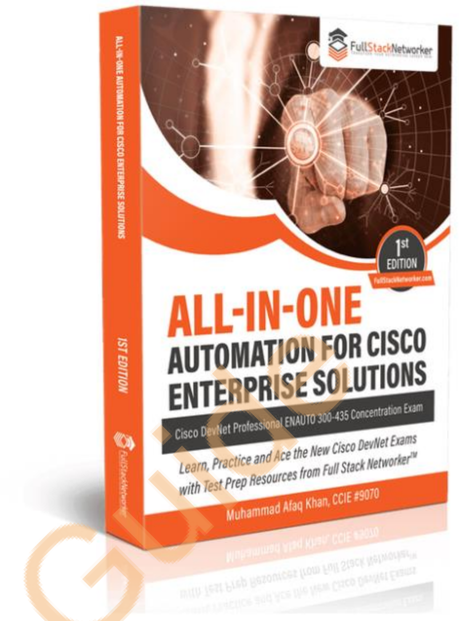
FullStackNetworker
1st EDITION
FullStackNetworker.com
ALL-IN-ONE
AUTOMATION FOR CISCO ENTERPRISE SOLUTIONS
Cisco DevNet Professional ENAUTO 300-435 Concentration Exam
Learn, Practice and Ace the New Cisco DevNet Exams with Test Prep Resources from Full Stack Networker™
Muhammad Afaq Khan, CCIE #9070
ALL-IN-ONE AUTOMATION FOR CISCO ENTERPRISE SOLUTIONS
1ST EDITION

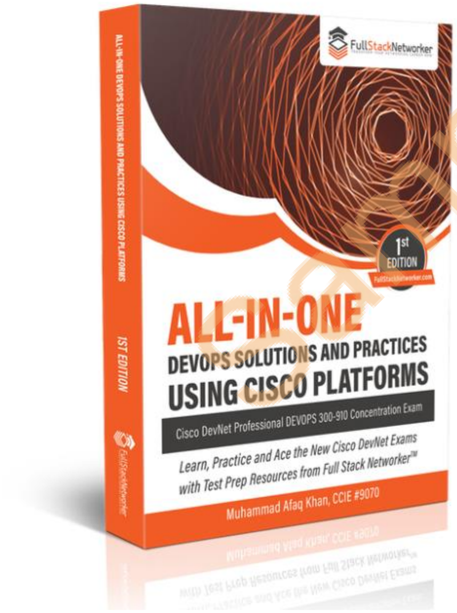
FullStackNetworker
1st EDITION
FullStackNetworker.com
ALL-IN-ONE
DEVOPS SOLUTIONS AND PRACTICES USING CISCO PLATFORMS
Cisco DevNet Professional DEVOPS 300-910 Concentration Exam
Learn, Practice and Ace the New Cisco DevNet Exams with Test Prep Resources from Full Stack Networker™
Muhammad Afaq Khan, CCIE #9070
ALL-IN-ONE DEVOPS SOLUTIONS AND PRACTICES USING CISCO PLATFORMS
1ST EDITION

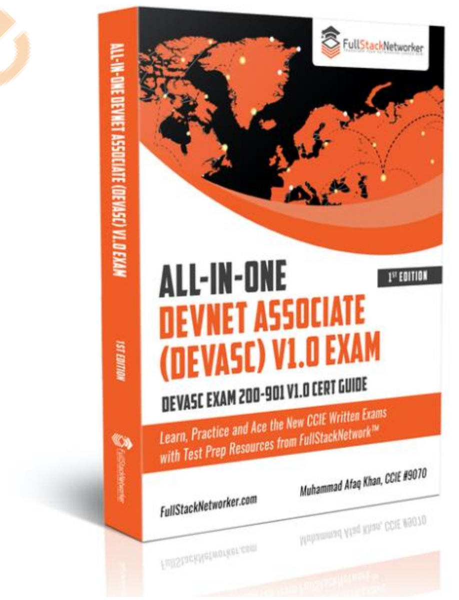
FullStackNetworker
ALL-IN-ONE
1st EDITION
DEVNET ASSOCIATE (DEVASC) V1.0 EXAM
DEVASC EXAM 200-901 V1.0 CERT GUIDE
Learn, Practice and Ace the New CCIE Written Exams with Test Prep Resources from FullStackNetwork™
FullStackNetworker.com
Muhammad Afaq Khan, CCIE #9070
ALL-IN-ONE DEVNET ASSOCIATE (DEVASC) V1.0 EXAM
1ST EDITION

Sample Guide

Our Cisco Next-Level Certification Catalog

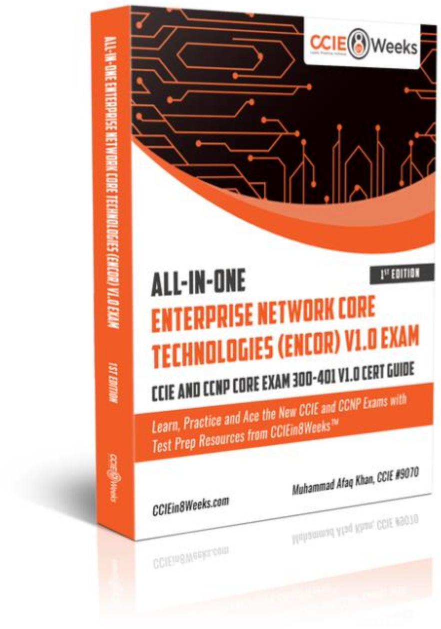
CCIE 8 Weeks
ALL-IN-ONE ENTERPRISE NETWORK CORE TECHNOLOGIES (ENCOR) V1.0 EXAM
1ST EDITION
ALL-IN-ONE
1ST EDITION
ENTERPRISE NETWORK CORE TECHNOLOGIES (ENCOR) V1.0 EXAM
CCIE AND CCNP CORE EXAM 300-401 V1.0 CERT GUIDE
Learn, Practice and Ace the New CCIE and CCNP Exams with Test Prep Resources from CCIEin8Weeks™
Muhammad Afaq Khan, CCIE #9070
CCIEin8Weeks.com

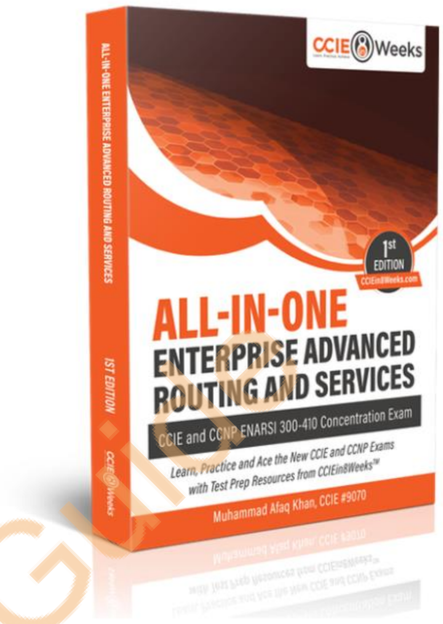
CCIE 8 Weeks
ALL-IN-ONE ENTERPRISE ADVANCED ROUTING AND SERVICES
1ST EDITION
1st EDITION
CCIEin8Weeks.com
ALL-IN-ONE
ENTERPRISE ADVANCED ROUTING AND SERVICES
CCIE and CCNP ENARSI 300-410 Concentration Exam
Learn, Practice and Ace the New CCIE and CCNP Exams with Test Prep Resources from CCIEin8Weeks™
Muhammad Afaq Khan, CCIE #9070

Sample Guide

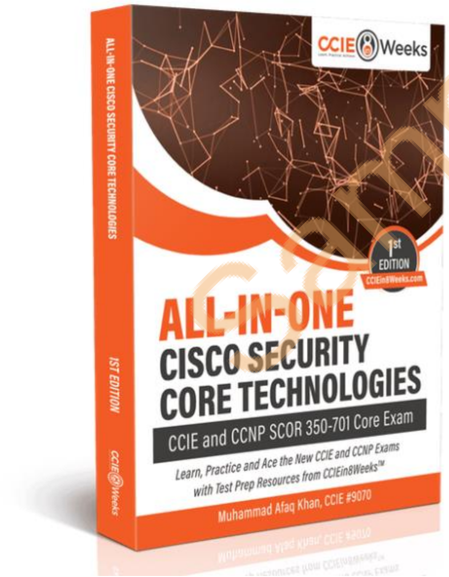
CCIE 8 Weeks
ALL-IN-ONE CISCO SECURITY CORE TECHNOLOGIES
1ST EDITION
1st EDITION
CCIEin8Weeks.com
ALL-IN-ONE
CISCO SECURITY CORE TECHNOLOGIES
CCIE and CCNP SCOR 350-701 Core Exam
Learn, Practice and Ace the New CCIE and CCNP Exams with Test Prep Resources from CCIEin8Weeks™
Muhammad Afaq Khan, CCIE #9070

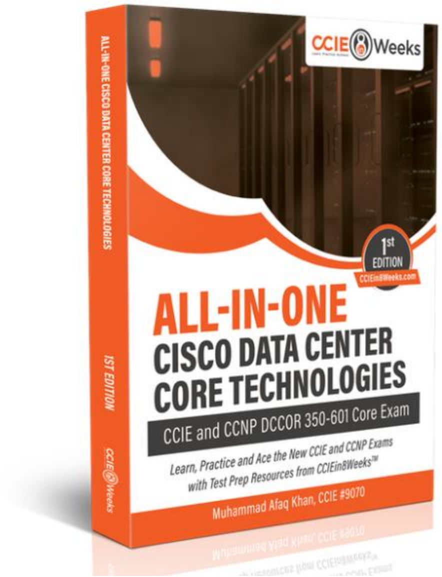
CCIE 8 Weeks
ALL-IN-ONE CISCO DATA CENTER CORE TECHNOLOGIES
1ST EDITION
1st EDITION
CCIEin8Weeks.com
ALL-IN-ONE
CISCO DATA CENTER CORE TECHNOLOGIES
CCIE and CCNP DCCOR 350-601 Core Exam
Learn, Practice and Ace the New CCIE and CCNP Exams with Test Prep Resources from CCIEin8Weeks™
Muhammad Afaq Khan, CCIE #9070

Contents at a Glance

Chapter 1 Network
Chapter 2 Compute
Chapter 3 Storage Network
Chapter 4 Automation
Chapter 5 Security

Sample Guide

Table of Contents

About the Author 11

Preface 12

What this Study Guide contains 13

How to use this Study Guide 13

What's available on the CCIEin8weeks website 14

CHAPTER 1 NETWORK 16

Network 17
Apply routing protocols 17
OSPFv2 and OSPFv3 17
Route Summarization 19
Inter-area Route Summarization 20
External Route Summarization 21
Route Filtering 22
OSPFv3 24
Further Reading 25
MP-BGP 25
PIM 25
Any Source Multicast (ASM) 26
Source Specific Multicast (SSM) 26
FHRP 26
FHRP Best Practices 28
Apply switching protocols such as RSTP +, LACP, and vPC 29
RSTP + / Rapid PVST+ 29
LACP 29
vPC 31
Apply overlay protocols such as VXLAN EVPN and OTV 31
VXLAN EVPN 32
VXLAN Tunnel Endpoint (VTEP) 33
Further Reading 34
OTV 34
Apply ACI concepts 34
Fabric setup 34
Access policies 35
VMM 36
Tenant policies 37
Analyze packet flow (unicast, multicast, and broadcast) 37
Analyze cloud service and deployment models (NIST 800-145) 38
Cloud Deployment Models 39
Cloud Service Models 39
Public Cloud 41
Private Cloud 42
Virtual Private Cloud (VPC) 43
Hybrid Cloud 43
Multi-cloud 44
Infrastructure as a service (IaaS) 46
Platform as a service (PaaS) 47
Software as a Service (SaaS) 47
Describe software updates and their impacts 49

Sample Guide

Disruptive 49
EPLD 49
Patches 49
Implement network configuration management 49
Implement infrastructure monitoring such as NetFlow and SPAN 53
NetFlow Interface Support for Ingress (Received) Traffic on an interface 54
NetFlow Interface Support for Egress (Transmitted) Traffic on an Interface 54
NetFlow Flow Export Destination and Version 54
Verifying NetFlow Configuration 54
Creating a Flow Record 57
Creating a Flow Monitor 57
Creating a Flow Exporter 57
Verifying Data in the Flow Monitor Cache 57
Configure and verify SPAN/RSPAN/ERSPAN 58
Configuring SPAN, RSPAN and ERSPAN 59
Configuring ERSPAN Source Session 59
Configuring ERSPAN Destination Session 60
Explain network assurance concepts such as streaming telemetry 60

Chapter Summary 62

CHAPTER 2 COMPUTE 64
Implement Cisco Unified Compute System Rack Servers 65
Implement Cisco Unified Compute System Blade Chassis 65
Initial setup 65
Infrastructure management 66
Network management (VLANs, pools and policies, templates, QoS) 66
Storage management (SAN connectivity, Fibre Channel zoning, VSANs, WWN pools, SAN policies, templates) 68
Server management (Server pools and boot policies) 69
Explain HyperFlex Infrastructure Concepts and benefits (Edge and Hybrid Architecture vs all-flash) 69
Describe firmware and software updates and their impacts on B-Series and C-Series servers 71
Implement compute configuration management (Backup and restore) 72
Implement infrastructure monitoring such as SPAN and Intersight 73
Further Reading 87

Chapter Summary 88

CHAPTER 3 STORAGE NETWORK 90
Implement Fibre Channel 91
Switch fabric initialization 91
Port channels 92
FCID 92
CFS 93
Further Reading 94
Zoning 94
FCNS 95
Device alias 95
Further Reading 96
NPV and NPIV 96
VSAN 97
Implement FCoE Unified Fabric (FIP and DCB) 97
Describe NFS and NAS concepts 101
Describe software updates and their impacts (Disruptive/nondisruptive and EPLD) 104
Implement infrastructure monitoring 105

Chapter Summary 106

Sample Guide

CHAPTER 4 AUTOMATION 108
Implement automation and scripting tools 109
EEM 109
Scheduler 110
Bash Shell and Guest Shell for NX-OS 111
REST API 120
Further Reading 123
Further Reading 131
JSON and XML encodings 131
XML Example 132
XML Prologue 133
XML Comments 133
XML Body 133
XML Attributes 134
XML Namespaces 134
JSON Example 134
JSON Data Types 135
JSON Objects 135
JSON Maps and Lists 135
YAML Example 135
YAML File Structure 136
YAML Data Types 136
YAML Indentation and File Structure 136
YAML Maps and Lists 136
YAML Comments 137
XML Parsing in Python 137
Using the ElementTree APIs to parse XML 138
Using Minidom Module to parse XML 140
JSON Parsing in Python 141
YAML Parsing in Python 142
Further Reading 143
Evaluate automation and orchestration technologies 144
Ansible and Puppet 144
Python 146
POAP 154
DCNM 155
UCSD 156
PowerShell 160
Chapter Summary 165
CHAPTER 5 SECURITY 167
Apply network security 168
AAA and RBAC 168
Configuring login authentication and authorization on Cisco IOS XE 168
Configuring downloadable ACLs (per-user ACLs) using Cisco IOS XE 168
ACI contracts and micro segmentation 169
First-hop security features such as dynamic ARP inspection (DAI), DHCP snooping, and port security .. 170
CoPP 172
Apply compute security 175
AAA and RBAC 175
Further Reading 176
Keychain authentication 176
Apply storage security 177
AAA and RBAC 177

Sample Guide

Port security 178
Fabric binding 178

Chapter Summary 180

Sample Guide

Sample Guide

About the Author

Muhammad Afaq Khan started his professional career at Cisco TAC San Jose and passed his first CCIE in 2002 (#9070). He held multiple technical and management positions at Cisco San Jose HQ over his 11 years of tenure at the company before moving into cloud software and data center infrastructure IT industries.

He has worked at startups as well as Fortune 100 companies in senior leadership positions over his career. He is also a published author (Cisco Press, 2009) and holds multiple patents in the areas of networking, security, and virtualization. Currently, he is a founder at Full Stack Networker and a vocal advocate for network automation technologies and NetDevOps.

Sample Guide

Preface

Congratulations! You have taken your first step towards preparing and passing the Cisco Data Center Core (DCCOR) 350-601 V1.0 Exam.

Did you just purchase a copy? Interested in getting access to an **DCCOR Exam Quiz**? Register here[1] and send us an email at support@cciein8weeks.com to get started.

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

Sample Guide

[1] https://bit.ly/2UhExXh

What this Study Guide contains

This study guide includes all the topics from Cisco's official exam blueprint for Implementing and Operating Cisco Data Center Core Technologies, i.e., DCCOR 350-601. As you may already have noticed on the "Contents at a Glance" page that this guide has been formatted around the Cisco's official DCCOR 350-601 exam topics or curriculum. The benefit? As you read through the various topics, you will know exactly where you're within your learning journey.

All contents are carefully covered with core concepts, configuration snippets (where applicable), and topic summaries to help you master the skills so you can confidently face the pressures of the Cisco DCCOR exam as well as its real-world application.

The DCCOR exam contains subject matter from six domains, which also happen to be the six chapters in this study guide.

- Network
- Compute
- Storage Network
- Automation
- Security

How to use this Study Guide

This guide is for anyone who's studying for Cisco DCCOR 350-601 V1.0 exam. I strongly suggest taking a methodical approach for exam preparation, i.e., start with a target date or when you would like to sit for the actual exam and then work backward to see what kind of study plan would work for you.

Cisco Data Center Core (EDCCOR) 350-601 V1.0 Exam Topics Bodies of Knowledge	Exam Weight
Network	25%
Compute	25%
Storage Network	20%
Automation	15%
Security	15%

Sample Guide

What’s available on the CCIEin8weeks website

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

- All-in-One Exam Prep Bundle[2] that covers all bodies of knowledge tested on the DCCOR Exam
- 6x Practice Quizzes (one for each section as per the official curriculum)
- 1x Practice Exam simulation (to help you prepare to face the pressure of a real Cisco exam)

Sample Guide

[2] https://bit.ly/3f7h6Xz

Sample Guide

CHAPTER 1 NETWORK

This chapter covers the following exam topics from Cisco's official 350-601 V1.0 Data Center Core Technologies (DCCOR)[3] exam blueprint.

- Network
- Apply routing protocols
  - OSPFv2 and OSPFv3
  - MP-BGP
  - PIM
  - FHRP
- Apply switching protocols such as RSTP +, LACP, and vPC
- Apply overlay protocols such as VXLAN, EVPN, and OTV
- Apply ACI concepts
  - Fabric setup
  - Access policies
  - VMM
  - Tenant policies
- Analyze packet flow (unicast, multicast, and broadcast)
- Analyze cloud service and deployment models (NIST 800-145)
- Describe software updates and their impacts
  - Disruptive
  - EPLD
  - Patches
- Implement network configuration management
- Implement infrastructure monitoring such as NetFlow and SPAN
- Explain network assurance concepts such as streaming telemetry

Sample Guide

[3] https://bit.ly/2YnmFdE

Network

Apply routing protocols

OSPFv2 and OSPFv3

Open Shortest Path First (OSPF) v2 protocol is defined in RFC 2328 and based on link-state technology where a link is an interface on an L3 device such as a router. The state of the link is an attribute of that interface and its relationship to its neighbors. The interface attributes relevant to OSPF include the IP address, subnet mask, the type of the network, the routers that are also connected to that network and what have you.

Unlike distance vector protocols that use Bellman-Ford, OSPF uses the shortest path first (or SPF or Dijkstra SPF) algorithm to determine the shortest to all known destination networks with the help of a graph.

When an OSPF router boots up, it generates a link-state advertisement (or LSA) for that router which represents the state of links. All routers exchange their LSAs via flooding mechanism. Once an exchange is completed, every router ends up with a database that is used to calculate the shortest path to each destination. Each OSPF router uses the Dijkstra SPF algorithm to derive the shortest path tree and the result of this calculation is stored in the routing table (or RIB). The algorithm places each router at the root of a tree and then calculates the shortest path to each destination network based on the cumulate cost needed to reach that destination.

	EIGRP	**OSPF**
Best Path Selection Algorithm	DUAL FSM	Dijkstra SPF
Administrative Distance	90	110
Metric	Bandwidth, Load, Delay and Reliability	Cost
VLSM	Supported	Not Supported
Authentication	Supported	Supported
Multi Path	Supported	Supported (ECMP)
Open Industry Standard	No	Yes, RFC 2328

Sample Guide

OSPF uses interface cost as its metric, which is inversely proportional to the bandwidth of that interface, i.e. higher bandwidth means lower cost by default unless you modify it using ip ospf cost <value> command.

OSPF uses flooding to exchange link-state advertisements between routers and all routers within an area have an exact link-state database. Routers that have interfaces in multiple areas including backbone are known as Area Border Routers (ABRs). Routers that act as a gateway between OSPF and other routing protocols or other instances of the OSPF process are known as Autonomous System Border Routers (or ASBRs). OSPF finite state machine (or FSM) includes eight different states including down, attempt, init, two-way, exstart, exchange, loading and full.

OSPF addresses three classes of network types, point to point (p2p), point to multipoint (p2mp) and broadcast.

Enabling OSPF on a Cisco router involves several steps.

1. Enabling an OSPF process (router ospf <process-id>)
2. Assigning areas to the interfaces (network <network/IP-address> <mask> <area-id>)
3. Configuring authentication (optional)

The OSPF process ID is a numeric value only locally significant to the router, i.e. it is never sent to the other routers and thus doesn’t have to match with process IDs on other routers either. Technically, you can also run multiple OSPF processes on the same router too.

The network command is a way of assigning an interface to an area whereas a mask is used for ease of configuration so that you can put a bunch of interfaces into an area with one line. Area-id is the area number you want the interface to be in. It can be configured in a simple number format such as 0 or 1 or 2, and in the form of an IP address say 0.0.0.0 or 1.1.1.1.

Sample Guide

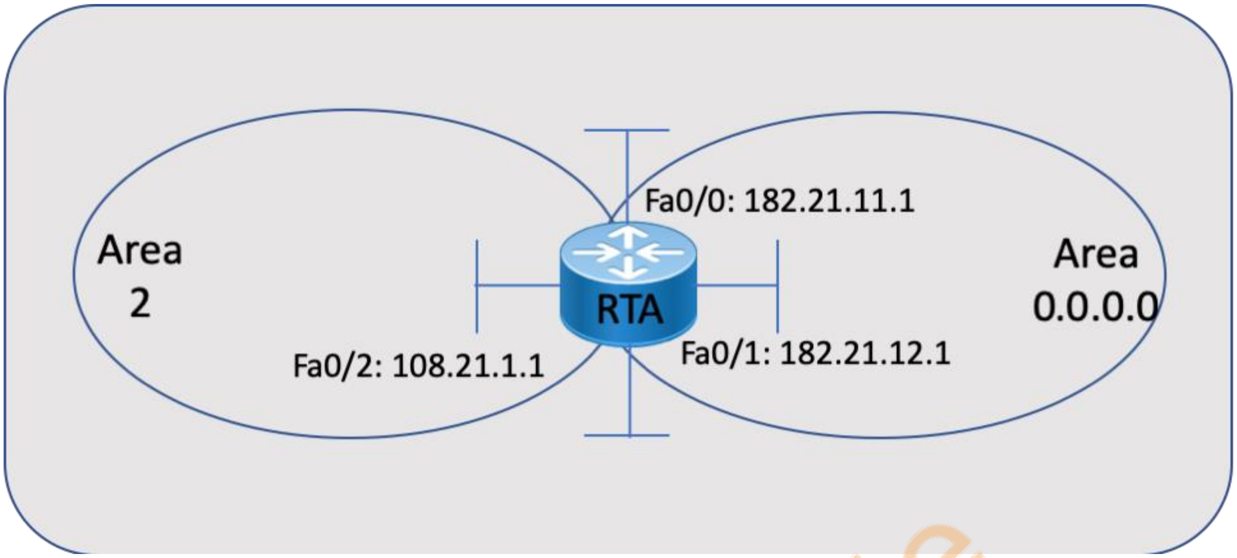

Sample Guide

```
RTA#
interface fa0/0
ip address 182.21.11.1 255.255.255.0

interface fa0/1
ip address 182.21.12.2 255.255.255.0

interface fa0/2
ip address 108.21.1.1 255.255.255.0

router ospf 101
network 182.21.0.0 0.0.255.255 area 0.0.0.0
network 108.21.1.1 0.0.0.0 area 2
```

Route Summarization

Summarization is about consolidating multiple routes into one single advertisement. In OSPF, this is normally done at the ABRs. You can configure summarization between any two areas; however, it is recommended to summarize towards the backbone area so it can inject those summaries into other areas. Summarization is highly effective if the network addresses assigned are contiguous.

There are two types of summarization, i.e.

- Inter-area routes
- External routes

Inter-area Route Summarization

Inter-area route summarization is done on ABRs and applies to routes from within the AS, i.e. it doesn't apply to routes coming into the OSPF domain from external sources. You can use area <area-id> range <address> <mask> command to configure inter-area summary. Here, area-id refers to the area containing networks that are to be summarized.

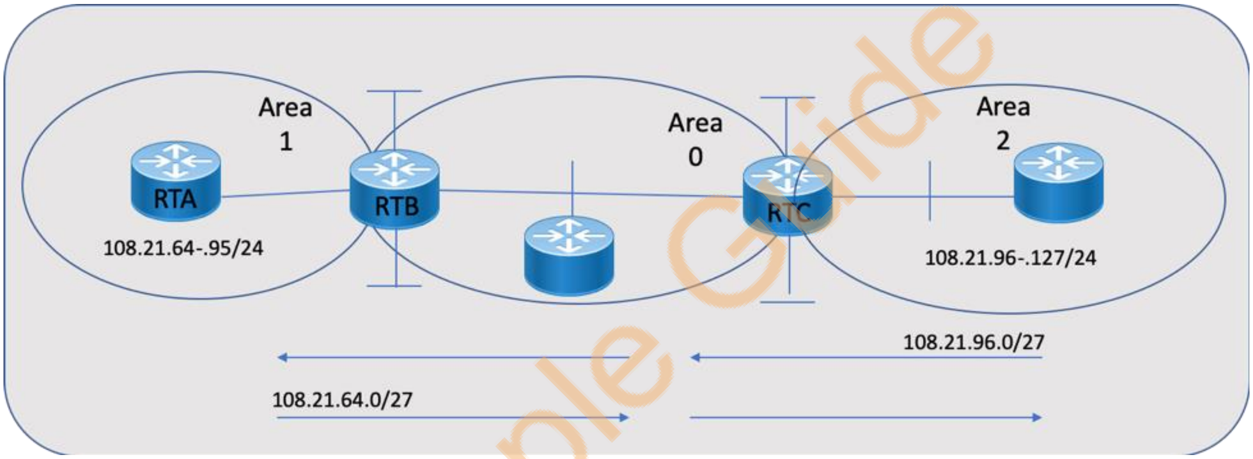

Sample Guide

In this topology, RTB is summarizing the range of subnets from 108.21.64.0 to 18.21.95.0 into one range, i.e. 108.21.64.0 255.255.224.0 (/27) into the backbone area. Likewise, RTC is summarizing 108.21.96.0/27 into the backbone.

```
RTB#
 router ospf 101
 area 1 range 108.21.64.0 255.255.224.0
```

```
RTC#
 router ospf 101
 area 1 range 108.21.96.0 255.255.224.0
```

External Route Summarization

External routes summarization is relevant to external routes only, ones that are injected into OSPF via redistribution. Much like inter-area summarization, the address range being contiguous would make it straightforward.

You will need to use summary-address <ip-address> <mask> command on ASBR(s) doing the redistribution into OSPF. This command will not affect if configured on a router with no connection to another router outside the OSPF domain.

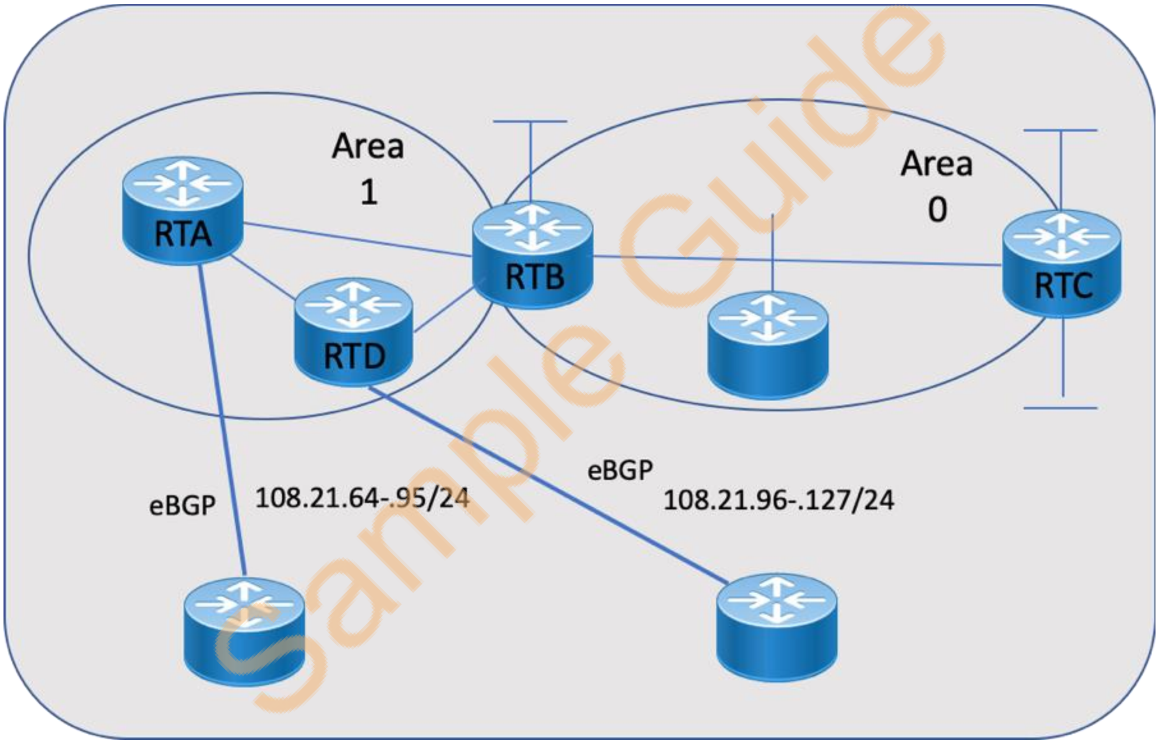

Sample Guide

In the above topology, RTA and RTD are injecting external routes into OSPF. RTA is injecting subnets within the range of 108.21.64-95 and RTD is injecting subnets 108.21.96-127.

```
RTA#
 router ospf 101
 summary-address 108.21.64.0 255.255.224.0
 redistribute bgp 50 metric 1000 subnets

 RTD#
 router ospf 101
```

```
summary-address 108.21.96.0 255.255.224.0
redistribute bgp 20 metric 1000 subnets
```

Route Filtering

Unlike RIP or distance vector protocols, OSPF has built-in controls over route propagation. OSPF routes are allowed or denied into different OSPF areas based on the area type, such as backbone or stub areas. OSPF ABRs limit the advertisement of different types of routes into different OSPF areas depending on the type of associated LSA. For example, an OSPF ABR bordering an OSPF stub area would prevent the advertisement of external routes into the stub area. The ABR is a stub or totally-stub area that would advertise a default route as an inter-area route. However, an ABR to a totally-stub area prevents advertisements of any inter-area including any external routes into that area. One common use of route filtering is when performing mutual redistribution.

There are a couple of key concepts to understand about OSPF filtering.

- With passive interface command configured, OSPF doesn’t send hellos on an interface. This means that the device wouldn’t discover any neighbor on that interface.
- Most filtering tools available do not filter out or remove routes from within the LS database.
- The route filters have no impact on the presence of routes in the routing table beyond the local router, i.e. they are only locally significant.

Distribute-list in works on any OSPF router and would prevent routes from being added to the routing table but routes still get added to the LS database i.e. the downstream neighbors will still have those routes. However, distribute-list out works on an ASBR to filter redistributed routes into other protocols.

Sample Guide

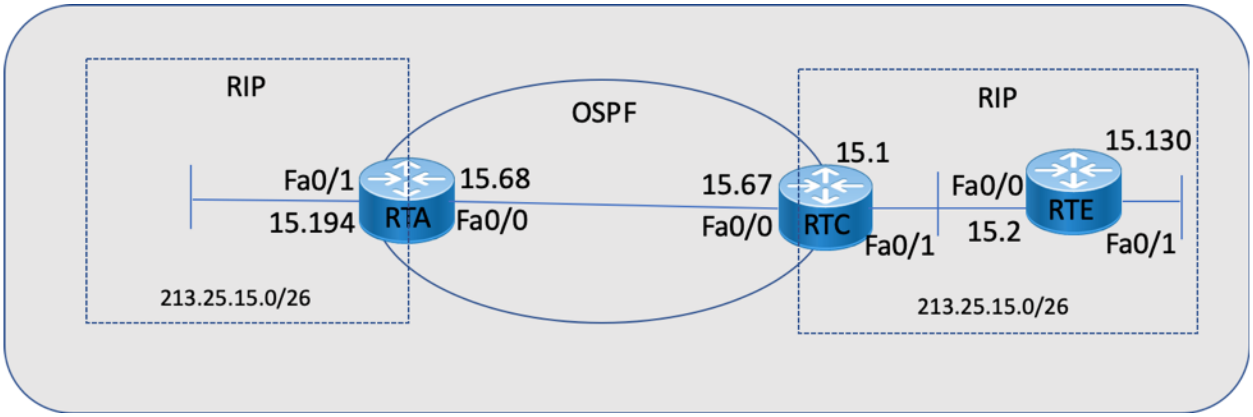

```
RTE#
 interface fa0/1
  ip address 213.25.15.130 255.255.255.192

 interface fa0/0
  ip address 213.25.15.2 255.255.255.192

 router rip
  network 213.25.15.0

RTC#
 interface fa0/0
  ip address 213.25.15.67 255.255.255.192

 interface fa0/1
  ip address 213.25.15.1 255.255.255.192

 router ospf 101
  redistribute rip metric 10 subnets
  network 213.25.15.0 0.0.0.255 area 0

 router rip
  redistribute ospf 101 metric 2
  passive-interface Ethernet0
  network 213.25.15.0

RTA#
interface fa0/0
 ip address 213.25.15.68 255.255.255.192

router ospf 101
 redistribute rip metric 10 subnets
 network 213.25.15.0 0.0.0.255 area 0

router rip
 redistribute ospf 101 metric 1
 network 213.25.15.0
```

If you were to do a “show ip route” on RTC, you would have found two paths to the 213.25.15.128 destination network. This occurred because RTC advertised the route to RTA via

Sample Guide

OSPF and RTA advertised it back via RIP. Now, in order to fix this issue, the most effective way would be to use a distribute-list on RTA to deny the 213.25.15.0 network from being put back into RIP.

```
RTA#
  interface fa0/0
  ip address 213.25.15.68 255.255.255.192

 router ospf 101
  redistribute rip metric 10 subnets
  network 213.25.15.0 0.0.0.255 area 0

 router rip
  redistribute ospf 101 metric 1
  network 213.25.15.0
  distribute-list 1 out ospf 101
```

OSPFv3

Open Shortest Path First (OSPF) v3 protocol is defined in RFC 5340 which defines the modifications to OSPFv2 to support IPv6. The fundamental ways the OSPF operates, such as flooding, DR election process, area support, SPF computations, NSSAs, and what have you, remain unchanged.

The OSPFv3 includes the following changes to support IPv6 protocol.

- Removal of addressing semantics from OSPF packets and the LSAs
- Introduction of new LSAs to carry IPv6 addresses and prefixes
- Ability to run OSPF on per-link basis as opposed to per-subnet basis
- Removal of authentication from the protocol (IPv6 already includes AH/ESP IPsec support)
- More flexibility around OSPF options handling, fields and packet size limitations
- Generalization of flooding scope for LSAs (link-local, area scope, and AS scope)

Sample Guide

Further Reading

OSPF Configuration Guide[4]
OSPF Design Guide[5]

Sample Guide

[4] https://bit.ly/2GMcKGh

[5] https://bit.ly/31jSOUx