



FortiOS™ Handbook
Advanced Routing for FortiOS 5.0



FortiOS™ Handbook - Advanced Routing for FortiOS 5.0

April 29, 2014

01-500-189263-20140429

Copyright© 2014 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, and FortiGuard®, are registered trademarks of Fortinet, Inc., and other Fortinet names herein may also be trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance metrics contained herein were attained in internal lab tests under ideal conditions, and performance may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to the performance metrics herein. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. Fortinet disclaims in full any guarantees. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.

Technical Documentation	docs.fortinet.com
Knowledge Base	kb.fortinet.com
Customer Service & Support	support.fortinet.com
Training Services	training.fortinet.com
FortiGuard	fortiguard.com
Document Feedback	techdocs@fortinet.com

Contents

Change Log	7
Introduction	8
Before you begin.....	8
How this guide is organized.....	8
Advanced Static Routing	9
Routing concepts.....	9
Routing in VDOMs	9
Default route	10
Adding a static route.....	10
Routing table.....	10
Building the routing table.....	17
Static routing security	17
Multipath routing and determining the best route	19
Route priority	20
Troubleshooting static routing	21
Static routing tips.....	23
Policy routing	24
Adding a policy route.....	25
Moving a policy route.....	27
Transparent mode static routing.....	27
Static routing example	28
Network layout and assumptions	28
General configuration steps.....	29
Get your ISP information such as DNS, gateway, etc.	30
Configure FortiGate unit.....	30
Configure Admin PC and Dentist PCs	35
Testing network configuration	36
Advanced static example: ECMP failover and load balancing	37
Equal-Cost Multi-Path (ECMP)	37
Configuring interface status detection for gateway load balancing	38
Configuring spillover or usage-based ECMP.....	40
Configuring weighted static route load balancing	42
Dynamic Routing Overview	44
What is dynamic routing?	44
Comparing static and dynamic routing.....	44
Dynamic routing protocols.....	45
Minimum configuration for dynamic routing	47

Comparison of dynamic routing protocols	47
Features of dynamic routing protocols	47
When to adopt dynamic routing	50
Choosing a routing protocol	52
Dynamic routing terminology	53
IPv6 in dynamic routing	58
Routing Information Protocol (RIP)	59
RIP background and concepts	59
Background.....	59
Parts and terminology of RIP	60
How RIP works	65
Troubleshooting RIP	70
Routing Loops.....	70
Holddowns and Triggers for updates	73
Split horizon and Poison reverse updates	73
Debugging IPv6 on RIPng.....	74
Simple RIP example.....	74
Network layout and assumptions	75
General configuration steps.....	76
Configuring the FortiGate units system information	76
Configuring FortiGate unit RIP router information	86
Configuring other networking devices	90
Testing network configuration	91
RIPng — RIP and IPv6.....	91
Network layout and assumptions	91
Configuring the FortiGate units system information	92
Configuring RIPng on FortiGate units.....	95
Configuring other network devices	96
Testing the configuration	96
Border Gateway Protocol (BGP)	98
BGP background and concepts	98
Background.....	98
Parts and terminology of BGP	98
How BGP works.....	107
Troubleshooting BGP	110
Clearing routing table entries.....	111
Route flap.....	111
Dual-homed BGP example	115
Network layout and assumptions	116
Configuring the FortiGate unit.....	118
Configuring other networking devices	127
Testing this configuration.....	127

Redistributing and blocking routes in BGP	129
Network layout and assumptions	129
Configuring the FortiGate unit.....	130
Testing network configuration	135
Open Shortest Path First (OSPF)	136
OSPF Background and concepts	136
Background.....	136
The parts and terminology of OSPF	136
How OSPF works.....	143
Troubleshooting OSPF	148
Clearing OSPF routes from the routing table.....	148
Checking the state of OSPF neighbors	149
Passive interface problems.....	149
Timer problems	149
Bi-directional Forwarding Detection (BFD)	149
Authentication issues.....	150
DR and BDR election issues.....	150
Basic OSPF example	150
Network layout and assumptions	151
Configuring the FortiGate units.....	152
Configuring OSPF on the FortiGate units	155
Configuring other networking devices	162
Testing network configuration	162
Advanced inter-area OSPF example	163
Network layout and assumptions	163
Configuring the FortiGate units.....	165
Configuring OSPF on the FortiGate units	169
Configuring other networking devices	173
Testing network configuration	173
Controlling redundant links by cost.....	173
Adjusting the route costs	175
Verifying route redundancy	177
Intermediate System to Intermediate System Protocol (IS-IS)	178
IS-IS background and concepts.....	178
Background.....	178
How IS-IS works	179
Parts and terminology of IS-IS.....	180
Troubleshooting IS-IS	185
Routing loops.....	185
Split horizon and Poison reverse updates	188

Simple IS-IS example	188
Network layout and assumptions	189
Expectations	189
CLI configuration.....	190
Verification	192
Troubleshooting	195
Index	196

Change Log

Date	Change Description
2014-04-29	Removed erroneous feature support.
2014-02-19	Included Advanced Static Routing section.
2014-02-01	New ISIS section.
2013-12-12	Updates to OSPF section.
2013-11-27	Updates to BGP section.
2013-11-06	Updates to BGP section.
2013-01-04	Initial Release.

Introduction

Dynamic routing is required in complex and changing network configurations where static routing does not provide sufficient convergence, redundancy, or other extended functionality.

This guide provides detailed information about FortiGate dynamic routing including common dynamic routing features, troubleshooting, and each of the protocols including RIP, BGP, and OSPF.

This chapter contains the following sections:

- [Before you begin](#)
- [How this guide is organized](#)

Before you begin

Before you begin using this guide, take a moment to note the following:

- This guide is based on the assumption that you are a FortiGate administrator.
- The configuration examples show steps for both the web-based manager (GUI) and the CLI. For more information about using the CLI, see the [FortiGate CLI Reference](#).
- At this stage, the following installation and configuration conditions are assumed:
 - You have administrative access to the web-based manager and CLI.

How this guide is organized

This chapter describes advanced static routing concepts and how to implement dynamic routing on FortiGate units.

This FortiOS Handbook chapter contains the following sections:

[Advanced Static Routing](#) explains universal and static routing concepts, equal cost multipath (ECMP) and load balancing, policy routing, and routing in transparent mode.

[Dynamic Routing Overview](#) provides an overview of dynamic routing, compares static and dynamic routing, and helps you decide which dynamic routing protocol is best for you.

[Routing Information Protocol \(RIP\)](#) describes a distance-vector routing protocol intended for small, relatively homogeneous networks.

[Border Gateway Protocol \(BGP\)](#) describes classless inter-domain routing, and aggregate routes. BGP is the only routing protocol to use TCP for a transport protocol.

[Open Shortest Path First \(OSPF\)](#) provides background on the specific protocol explaining terms used and how the protocol works, as well as providing some troubleshooting information and examples on configuring the protocols in different situations.

[Intermediate System to Intermediate System Protocol \(IS-IS\)](#), which describes the link state protocol, is well-suited to smaller networks and with near universal support on routing hardware. The section also provides troubleshooting information and configuration examples.

Advanced Static Routing

Advanced static routing includes features and concepts that are used in more complex networks. Dynamic routing is not addressed in this section.

This section includes:

- Routing concepts
- Static routing tips
- Policy routing
- Transparent mode static routing
- Static routing example
- Advanced static example: ECMP failover and load balancing

Routing concepts

Many routing concepts apply to static routing. However without first understanding these basic concepts, it is difficult to understand the more complex dynamic routing.

This section includes:

- Routing in VDOMs
- Default route
- Adding a static route
- Routing table
- Building the routing table
- Static routing security
- Multipath routing and determining the best route

Routing in VDOMs

Routing on FortiGate units is configured per-VDOM. This means if VDOMs are enabled, you must enter a VDOM to do any routing configuration. This allows each VDOM to operate independently, with its own default routes and routing configuration.

In this guide, the procedures assume your FortiGate unit has VDOMs disabled. This is stated in the assumptions for the examples. If you have VDOMs enabled you will need to perform the following steps in addition to the procedure's steps.

To route in VDOMs - web-based manager

Select the VDOM that you want to view or configure at the bottom of the main menu.

To route in VDOMs - CLI

Before following any CLI routing procedures with VDOMs enabled, enter the following commands. For this example, it is assumed you will be working in the root VDOM. Change root to the name of your selected VDOM as needed.

```
config vdom
edit root
```

Following these commands, you can enter any routing CLI commands as normal.

Default route

The default route is used if either there are no other routes in the routing table or if none of the other routes apply to a destination. Including the gateway in the default route gives all traffic a next-hop address to use when leaving the local network. The gateway address is normally another router on the edge of the local network.

All routers, including FortiGate units, are shipped with default routes in place. This allows customers to set up and become operational more quickly. Beginner administrators can use the default route settings until a more advanced configuration is warranted.

FortiGate units come with a default static route with an IPv4 address of 0.0.0.0, an administration distance of 10, and a gateway IPv4 address.

Adding a static route

To add or edit a static route, go to *Router > Static > Static Routes* and select *Create New*.

Destination IP / Mask	Enter the destination IP address and netmask. A value of 0.0.0.0/0.0.0.0 is universal.
Device	Select the name of the interface which the static route will connect through.
Gateway	Enter the gateway IP address.
Distance	Enter the distance value, which will affect which routes are selected first by different protocols for route management or load balancing. The default is 10.
Priority	Enter the priority if desired, which will artificially weight the route during route selection. The higher the number, the less likely the route is to be selected over others. The default is 0.

Routing table

When two computers are directly connected, there is no need for routing because each computer knows exactly where to find the other computer. They communicate directly.

Networking computers allows many computers to communicate with each other. This requires each computer to have an IP address to identify its location to the other computers. This is much like a mailing address - you will not receive your postal mail at home if you do not have an address for people to send mail to. The routing table on a computer is much like an address book used to mail letters to people in that the routing table maintains a list of how to reach computers. Routing tables may also include information about the quality of service (QoS) of the route, and the interface associated with the route if the device has multiple interfaces.

Looking at routing as delivering letters is more simple than reality. In reality, routers loose power or have bad cabling, network equipment is moved without warning, and other such events happen that prevent static routes from reaching their destinations. When any changes such as these happen along a static route, traffic can no longer reach the destination — the route goes down. Dynamic routing can address these changes to ensure traffic still reaches its destination. The process of realizing there is a problem, backtracking and finding a route that is operational is called convergence. If there is fast convergence in a network, users won't even know that re-routing is taking place.

The routing table for any device on the network has a limited size. For this reason, routes that aren't used are replaced by new routes. This method ensures the routing table is always populated with the most current and most used routes—the routes that have the best chance of being reused. Another method used to maintain the routing table's size is if a route in the table and a new route are to the same destination, one of the routes is selected as the best route to that destination and the other route is discarded.

Routing tables are also used in unicast reverse path forwarding (uRPF). In uRPF, the router not only looks up the destination information, but also the source information to ensure that it exists. If there is no source to be found, then that packet is dropped because the router assumes it to be an error or an attack on the network.

The routing table is used to store routes that are learned. The routing table for any device on the network has a limited size. For this reason, routes that aren't used are replaced by new routes. This method ensures the routing table is always populated with the most current and most used routes — the routes that have the best chance of being reused. Another method used to maintain the routing table's size is if a route in the table and a new route are to the same destination, one of the routes is selected as the best route to that destination and the other route is discarded.

Some actions you can perform on the routing table include:

- Viewing the routing table in the web-based manager
- Viewing the routing table in the CLI
- Searching the routing table

Viewing the routing table in the web-based manager

VDOM

By default, all routes are displayed in the Routing Monitor list. The default static route is defined as 0.0.0.0/0, which matches the destination IP address of “any/all” packets.

To display the routes in the routing table, go to *Router > Monitor > Routing Monitor*.

Figure 1 shows the Routing Monitor list belonging to a FortiGate unit that has interfaces named “port1”, “port4”, and “lan”. The names of the interfaces on your FortiGate unit may be different.

Figure 2 shows the Routing Monitor list when IPv6 has been selected. Note that the information available for IPv6 is limited.

Figure 1: Routing Monitor list - IPv4

IP Version: IPv4		Type: All	Network:	Gateway:	Apply Filter		
Type	Subtype	Network	Distance	Metric	Gateway	Interface	Up Time (d h:m:s)
Connected		10.10.10.0/24	0	0	0.0.0.0	port4	
Connected		172.20.120.0/24	0	0	0.0.0.0	port1	

Figure 2: Routing Monitor list - IPv6

IP Version: IPv6		Interface	Network	Gateway
		havdlink1	fe80::10	
		havdlink1	ff00::8	

IP version:	Select IPv4 or IPv6. This is available only when IPv6 is enabled in the web-based manager. The fields displayed in the table depend on which IP version is selected.
Type:	<p>Select one of the following route types to search the routing table and display routes of the selected type only:</p> <p><i>All</i> — all routes recorded in the routing table.</p> <p><i>Connected</i> — all routes associated with direct connections to FortiGate unit interfaces.</p> <p><i>Static</i> — the static routes that have been added to the routing table manually.</p> <p><i>RIP</i> — all routes learned through RIP. For more information see “Routing Information Protocol (RIP)” on page 59.</p> <p><i>RIPNG</i> — all routes learned through RIP version 6 (which enables the sharing of routes through IPv6 networks).</p> <p><i>BGP</i> — all routes learned through BGP. For more information see “Border Gateway Protocol (BGP)” on page 98.</p> <p><i>OSPF</i> — all routes learned through OSPF. For more information see “Open Shortest Path First (OSPF)” on page 136.</p> <p><i>OSPF6</i> — all routes learned through OSPF version 6 (which enables the sharing of routes through IPv6 networks).</p> <p><i>IS-IS</i> — all routes learned through IS-IS. For more information see “Intermediate System to Intermediate System Protocol (IS-IS)” on page 178.</p> <p><i>HA</i> — RIP, OSPF, and BGP routes synchronized between the primary unit and the subordinate units of a high availability (HA) cluster. HA routes are maintained on subordinate units and are visible only if you are viewing the router monitor from a virtual domain that is configured as a subordinate virtual domain in a virtual cluster.</p> <p>Not displayed when IP version IPv6 is selected.</p> <p>For details about HA routing synchronization, see the FortiGate HA User Guide.</p>
Network:	Enter an IP address and netmask (for example, 172.16.14.0/24) to search the routing table and display routes that match the specified network.
Gateway:	Enter an IP address and netmask (for example, 192.168.12.1/32) to search the routing table and display routes that match the specified gateway.
Apply Filter	<p>Select to search the entries in the routing table based on the specified search criteria and display any matching routes.</p> <p>Not displayed when IP version IPv6 is selected.</p>
Type	<p>The type values assigned to FortiGate unit routes (Static, Connected, RIP, OSPF, or BGP).</p> <p>Not displayed when IP version IPv6 is selected.</p>

Subtype	<p>If applicable, the subtype classification assigned to OSPF routes.</p> <p>An empty string implies an intra-area route. The destination is in an area to which the FortiGate unit is connected.</p> <p><i>OSPF inter area</i> — the destination is in the OSPF AS, but the FortiGate unit is not connected to that area.</p> <p><i>External 1</i> — the destination is outside the OSPF AS. This is known as OSPF E1 type. The metric of a redistributed route is calculated by adding the external cost and the OSPF cost together.</p> <p><i>External 2</i> — the destination is outside the OSPF AS. This is known as OSPF E2 type. In this case, the metric of the redistributed route is equivalent to the external cost only, expressed as an OSPF cost.</p> <p><i>OSPF NSSA 1</i> — same as External 1, but the route was received through a not-so-stubby area (NSSA).</p> <p><i>OSPF NSSA 2</i> — same as External 2, but the route was received through a not-so-stubby area.</p> <p>For more information on OSPF subtypes, see “OSPF Background and concepts” on page 136.</p> <p>Not displayed when IP version 6 is selected.</p>
Network	<p>The IP addresses and network masks of destination networks that the FortiGate unit can reach.</p>
Distance	<p>The administrative distance associated with the route. A value of 0 means the route is preferable compared to other routes to the same destination, and the FortiGate unit may routinely use the route to communicate with neighboring routers and access servers.</p> <p>Modifying this distance for dynamic routes is route distribution. See “Redistributing and blocking routes in BGP” on page 129</p> <p>Not displayed when IP version 6 is selected.</p>
Metric	<p>The metric associated with the route type. The metric of a route influences how the FortiGate unit dynamically adds it to the routing table. The following are types of metrics and the protocols they are applied to.</p> <p><i>Hop count</i> — routes learned through RIP.</p> <p><i>Relative cost</i> — routes learned through OSPF.</p> <p><i>Multi-Exit Discriminator (MED)</i> — routes learned through BGP. However, several attributes in addition to MED determine the best path to a destination network. For more information on BGP attributes, see “BGP attributes” on page 104. By default, the MED value associated with a BGP route is zero. However, the MED value can be modified dynamically. If the value was changed from the default, the Metric column will display a non-zero value.</p> <p>Not displayed when IP version 6 is selected.</p>
Gateway	<p>The IP addresses of gateways to the destination networks.</p>

Interface	The interface through which packets are forwarded to the gateway of the destination network.
Up Time	The total accumulated amount of time that a route learned through RIP, OSPF, or BGP has been reachable. Not displayed when IP version IPv6 is selected.

Viewing the routing table in the CLI

In the CLI, you can easily view the static routing table just as in the web-based manager or you can view the full routing table.

When viewing the list of static routes using the CLI command `get route static`, it is the configured static routes that are displayed. When viewing the routing table using the CLI command `get router info routing-table all`, it is the entire routing table information that is displayed including configured and learned routes of all types. The two are different information in different formats.



If VDOMs are enabled on your FortiGate unit, all routing related CLI commands must be performed within a VDOM and not in the global context.

To view the routing table

```
# get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
    area
* - candidate default

S*    0.0.0.0/0 [10/0] via 192.168.183.254, port2
S     1.0.0.0/8 [10/0] via 192.168.183.254, port2
S     2.0.0.0/8 [10/0] via 192.168.183.254, port2
C     10.142.0.0/23 is directly connected, port3
B     10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
C     192.168.182.0/23 is directly connected, port2
```

Examining an entry:

```
B     10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
```

B	BGP. The routing protocol used.
10.160.0.0/23	The destination of this route including netmask.
[20/0]	20 indicates administrative distance of 20 out of a range of 0 to 255. 0 is an additional metric associated with this route, such as in OSPF
10.142.0.74	The gateway, or next hop.
port3	The interface used by this route.
2d18h02m	How old this route is, in this case almost three days old.

To view the kernel routing table

```
# get router info kernel
```

```
tab=254 vf=0 scope=253 type=1 proto=2 prio=0
  0.0.0.0/0.0.0.0/0->10.11.201.0/24 pref=10.11.201.4 gwy=0.0.0.0
  dev=5(external1)
```

```
tab=254 vf=0 scope=253 type=1 proto=2 prio=0
  0.0.0.0/0.0.0.0/0->172.20.120.0/24 pref=172.20.120.146 gwy=0.0.0.0
  dev=6(internal)
```

The parts of the routing table entry are:

tab	table number. This will be either 254 (unicast) or 255 (multicast).
vf	virtual domain of the firewall. This is the vdom index number. If vdoms are not enabled, this number will be 0.
type	type of routing connection. Valid values include: 0 - unspecified 1 - unicast 2 - local 3 - broadcast 4 - anycast 5 - multicast 6 - blackhole 7 - unreachable 8 - prohibited

proto	type of installation. This indicates where the route came from. Valid values include: 0 - unspecified 2 - kernel 11 - ZebOS routing module 14 - FortiOS 15 - HA 16 - authentication based 17 - HA1
prio	priority of the route. Lower priorities are preferred.
->10.11.201.0/24 (->x.x.x.x/mask)	the IP address and subnet mask of the destination
pref	preferred next hop along this route
gwy	gateway - the address of the gateway this route will use
dev	outgoing interface index. This number is associated with the interface for this route, and if VDOMs are enabled the VDOM will be included here as well. If an interface alias is set for this interface it will also be displayed here.

Searching the routing table

You can apply a filter to search the routing table and display certain routes only. For example, you can display one or more static routes, connected routes, routes learned through RIP, OSPF, or BGP, and routes associated with the network or gateway that you specify.

If you want to search the routing table by route type and further limit the display according to network or gateway, all of the values that you specify as search criteria must match corresponding values in the same routing table entry in order for that entry to be displayed — an implicit AND condition is applied to all of the search parameters you specify.

For example, if the FortiGate unit is connected to network 172.16.14.0/24 and you want to display all directly connected routes to network 172.16.14.0/24, you must select *Connected* from the *Type* list, type 172.16.14.0/24 in the *Network* field, and then select *Apply Filter* to display the associated routing table entry or entries. Any entry that contains the word “Connected” in its *Type* field and the specified value in the *Gateway* field will be displayed.

In this example, you will apply a filter to search for an entry for static route to 10.10.10.10/24

To search the FortiGate unit routing table in the web-based manager

1. Go to *Router > Monitor > Routing Monitor*.
2. From the *Type* list, select the type of route to display. In our example, select *Static*.
3. If you want to display routes to a specific network, type the IP address and netmask of the network in the *Networks* field. In our example, enter 10.10.10.10/24.
4. If you want to display routes to a specific gateway, type the IP address of the gateway in the *Gateway* field.

5. Select *Apply Filter*.



All of the values that you specify as search criteria must match corresponding values in the same routing table entry in order for that entry to be displayed.

To search the FortiGate unit routing table in the CLI

```
FGT # get router info routing-table details 10.10.10.10
Routing entry for 10.10.10.10/24
    Known via "static", distance 10, metric 0, best
```

If there are multiple routes that match your filter, they will all be listed, with the best match at the top of the list as indicated by the word best.

Building the routing table

In the factory default configuration, the FortiGate unit routing table contains a single static default route. You can add routing information to the routing table by defining additional static routes.

It is possible that the routing table is faced with several different routes to the same destination—the IP addresses of the next-hop router specified in those routes or the FortiGate interfaces associated with those routes may vary. In this situation, the “best” route is selected from the table.

The FortiGate unit selects the “best” route for a packet by evaluating the information in the routing table. The “best” route to a destination is typically associated with the shortest distance between the FortiGate unit and the closest gateway, also known as a next-hop router. In some cases, the next best route may be selected if the best route is unavailable.

The FortiGate unit installs the best available routes in the unit’s forwarding table, which is a subset of the unit’s routing table. Packets are forwarded according to the information in the forwarding table.

Static routing security

Securing the information on your company network is a top priority for network administrators. Security is also required as the routing protocols used are internationally known standards that typically provide little or no inherent security by themselves.

The two reasons for securing your network are the sensitive and proprietary information on your network, and also your external bandwidth. Hackers not only can steal your information, but they can also steal your bandwidth. Routing is a good low level way to secure your network, even before UTM features are applied.

Routing provides security to your network in a number of ways including obscuring internal network addresses with NAT and blackhole routing, using RPF to validate traffic sources, and maintaining an access control list (ACL) to limit access to the network.

This section includes:

- [Network Address Translation \(NAT\)](#)
- [Access Control List \(ACL\)](#)
- [Blackhole Route](#)
- [Reverse path lookup](#)

Network Address Translation (NAT)

Network address translation (NAT) is a method of changing the address traffic appears to originate from. This practice is used to hide the IP address on company's internal networks, and helps prevent malicious attacks that use those specific addresses.

This is accomplished by the router connected to that local network changing all the IP addresses to its externally connected IP address before sending the traffic out to the other networks, such as the Internet. Incoming traffic uses the established sessions to determine which traffic goes to which internal IP address. This also has the benefit of requiring only the router to be very secure against external attacks, instead of the whole internal network as would be the case without NAT. Securing one computer is much cheaper and easier to maintain.

Configuring NAT on your FortiGate unit includes the following steps.

1. Configure your internal network. For example use the `10.11.101.0` subnet.
2. Connect your internal subnet to an interface on your FortiGate unit. For example use `port1`.
3. Connect your external connection, for example an ISP gateway of `172.20.120.2`, to another interface on your Fortigate unit, for example `port2`.
4. Configure security policies to allow traffic between `port1` and `port2` on your FortiGate unit, ensuring that the NAT feature is enabled.

The above steps show that traffic from your internal network will originate on the `10.11.101.0` subnet and pass on to the `172.20.120.0` network. The FortiGate unit moves the traffic to the proper subnet. In doing that, the traffic appears to originate from the FortiGate unit interface on that subnet — it does not appear to originate from where it actually came from.

NAT “hides” the internal network from the external network. This provides security through obscurity. If a hacker tries to directly access your network, they will find the Fortigate unit, but will not know about your internal network. The hacker would have to get past the security-hardened FortiGate unit to gain access to your internal network. NAT will not prevent hacking attempts that piggy back on valid connections between the internal network and the outside world. However other UTM security measures can deal with these attempts.

Another security aspect of NAT is that many programs and services have problems with NAT. Consider if someone on the Internet tries to initiate a chat with someone on the internal network. The outsider only can access the FortiGate unit's external interface unless the security policy allows the traffic through to the internal network. If allowed in, the proper internal user would respond to the chat. However if its not allowed, the request to chat will be refused or time-out. This is accomplished in the security policy by allowing or denying different protocols.

Access Control List (ACL)

An access control list (ACL) is a table of addresses that have permission to send and receive data over a router's interface or interfaces. The router maintains an ACL, and when traffic comes in on a particular interface it is buffered, while the router looks up in the ACL if that traffic is allowed over that port or not. If it is allowed on that incoming interface, then the next step is to check the ACL for the destination interface. If the traffic passes that check as well the buffered traffic is delivered to its accentuation. If either of those steps fail the ACL check, the traffic is dropped and an error message may be sent to the sender. The ACL ensures that traffic follows expected paths, and any unexpected traffic is not delivered. This stops many network attacks. However, to be effective the ACL must be kept up to date — when employees or computers are removed from the internal network their IP addresses must also be removed from the ACL. For more information on the ACL, see the router chapter of the [FortiGate CLI Reference](#).

Blackhole Route

A blackhole route is a route that drops all traffic sent to it. It is very much like `/dev/null` in Linux programming.

Blackhole routes are used to dispose of packets instead of responding to suspicious inquiries. This provides added security since the originator will not discover any information from the target network.

Blackhole routes can also limit traffic on a subnet. If some subnet addresses are not in use, traffic to those addresses (traffic which may be valid or malicious) can be directed to a blackhole for added security and to reduce traffic on the subnet.

The loopback interface, a virtual interface that does not forward traffic, was added to enable easier configuration of blackhole routing. Similar to a normal interface, this loopback interface has fewer parameters to configure, and all traffic sent to it stops there. Since it cannot have hardware connection or link status problems, it is always available, making it useful for other dynamic routing roles. Once configured, you can use a loopback interface in security policies, routing, and other places that refer to interfaces. You configure this feature only from the CLI. For more information, see the system chapter of the *FortiGate CLI Reference*.

Reverse path lookup

Whenever a packet arrives at one of the FortiGate unit's interfaces, the unit determines whether the packet was received on a legitimate interface by doing a reverse lookup using the source IP address in the packet header. This is also called anti-spoofing. If the FortiGate unit cannot communicate with the computer at the source IP address through the interface on which the packet was received, the FortiGate unit drops the packet as it is likely a hacking attempt.

If the destination address can be matched to a local address (and the local configuration permits delivery), the FortiGate unit delivers the packet to the local network. If the packet is destined for another network, the FortiGate unit forwards the packet to a next-hop router according to a policy route and the information stored in the FortiGate forwarding table.

Multipath routing and determining the best route

Multipath routing occurs when more than one entry to the same destination is present in the routing table. When multipath routing happens, the FortiGate unit may have several possible destinations for an incoming packet, forcing the FortiGate unit to decide which next-hop is the best one.

It should be noted that some IP addresses will be rejected by routing protocols. These are called Martian addresses. They are typically IP addresses that are invalid and not routable because they have been assigned an address by a misconfigured system, or are spoofed addresses.

Two methods to manually resolve multiple routes to the same destination are to lower the administrative distance of one route or to set the priority of both routes. For the FortiGate unit to select a primary (preferred) route, manually lower the administrative distance associated with one of the possible routes. Setting the priority on the routes is a FortiGate unit feature and may not be supported by non-Fortinet routers.

Administrative distance is based on the expected reliability of a given route. It is determined through a combination of the number of hops from the source and the protocol used. A hop is when traffic moves from one router to the next. More hops from the source means more possible points of failure. The administrative distance can be from 1 to 255, with lower numbers being preferred. A distance of 255 is seen as infinite and will not be installed in the routing table.

Here is an example to illustrate how administrative distance works — if there are two possible routes traffic can take between two destinations with administrative distances of 5 (always up) and 31 (sometimes not available), the traffic will use the route with an administrative distance of 5. If for some reasons the preferred route (admin distance of 5) is not available, the other route will be used as a backup.

Different routing protocols have different default administrative distances. These different administrative distances are based on a number of factors of each protocol such as reliability,

speed, and so on. The default administrative distances for any of these routing protocols are configurable.

Table 1: Default administrative distances for routing protocols and connections

Routing protocol	Default administrative distance
Direct physical connection	1
Static	10
EBGP	20
OSPF	110
IS-IS	115
RIP	120
IBGP	200

Another method to determine the best route is to manually change the priority of both routes in question. If the next-hop administrative distances of two routes on the FortiGate unit are equal, it may not be clear which route the packet will take. Manually configuring the priority for each of those routes will make it clear which next-hop will be used in the case of a tie. The priority for a route be set in the CLI, or when editing a specific static route, as described in the next section. Lower priority routes are preferred. Priority is a Fortinet value that may or may not be present in other brands of routers.

All entries in the routing table are associated with an administrative distance. If the routing table contains several entries that point to the same destination (the entries may have different gateways or interface associations), the FortiGate unit compares the administrative distances of those entries first, selects the entries having the lowest distances, and installs them as routes in the FortiGate unit forwarding table. As a result, the FortiGate unit forwarding table contains only those routes having the lowest distances to every possible destination. While only static routing uses administrative distance as its routing metric, other routing protocols such as RIP can use metrics that are similar to administrative distance.

Route priority

After the FortiGate unit selects static routes for the forwarding table based on their administrative distances, the priority field of those routes determines routing preference. Priority is a Fortinet value that may or may not be present in other brands of routers.

You can configure the priority field through the CLI or the web-based manager. Priority values can range from 0 to 4 294 967 295. The route with the lowest value in the priority field is considered the best route. It is also the primary route.

To change the priority of a route - web-based manager

1. Go to *Router > Static > Static Routes*.
2. Select the route entry, and select *Edit*.
3. Select *Advanced*.
4. Enter the *Priority* value.
5. Select *OK*.

To change the priority of a route - CLI

The following command changes the priority to 5 for a route to the address 10.10.10.1 on the port1 interface.

```
config router static
  edit 1
    set device port1
    set gateway 10.10.10.10
    set dst 10.10.10.1
    set priority 5
  end
```

If there are other routes set to priority 10, the route set to priority 5 will be preferred. If there are routes set to priorities less than 5, those other routes will be preferred instead.

In summary, because you can use the CLI to specify which sequence numbers or priority field settings to use when defining static routes, you can prioritize routes to the same destination according to their priority field settings. For a static route to be the preferred route, you must create the route using the `config router static` CLI command and specify a low priority for the route. If two routes have the same administrative distance and the same priority, then they are equal cost multipath (ECMP) routes.

Since this means there is more than one route to the same destination, it can be confusing which route or routes to install and use. However, if you have enabled load balancing with ECMP routes, then different sessions will resolve this problem by using different routes to the same address.

Troubleshooting static routing

When there are problems with your network that you believe to be static routing related, there are a few basic tools available to locate the problem.

These tools include:

- [Ping](#)
- [Traceroute](#)
- [Examine routing table contents](#)

Ping

Beyond the basic connectivity information, ping can tell you the amount of packet loss (if any), how long it takes the packet to make the round trip, and the variation in that time from packet to packet.

If there is no packet loss detected, your basic network connectivity is OK.

If there is some packet loss detected, you should investigate:

- possible ECMP, split horizon, network loops
- cabling to ensure no loose connections

If there is total packet loss, you should investigate:

- hardware - ensure cabling is correct, and all equipment between the two locations is accounted for
- addresses and routes - ensure all IP addresses and routing information along the route is configured as expected
- firewalls - ensure all firewalls are set to allow PING to pass through

To ping from a Windows PC

1. Go to a DOS prompt. Typically you go to *Start > Run*, enter `cmd` and select OK.
2. Enter `ping 10.11.101.100` to ping the default internal interface of the FortiGate unit with four packets.

To ping from an Apple computer

1. Open the Terminal.
2. Enter `ping 10.11.101.100`.
3. If the ping fails, it will stop after a set number of attempts. If it succeeds, it will continue to ping repeatedly. Press `Control+C` to end the attempt and see gathered data.

To ping from a Linux PC

1. Go to a command line prompt.
2. Enter `"/bin/etc/ping 10.11.101.101"`.

Traceroute

Where ping will only tell you if it reached its destination and came back successfully, traceroute will show each step of its journey to its destination and how long each step takes. If ping finds an outage between two points, traceroute can be used to locate exactly where the problem is.

To use traceroute on a Windows PC

1. Go to a DOS prompt. Typically you go to *Start > Run*, enter `"cmd"` and select OK.
2. Enter `"tracert fortinet.com"` to trace the route from the PC to the Fortinet website.

To use traceroute from an Apple computer

1. Open the Terminal.
2. Enter `traceroute fortinet.com`.
3. The terminal will list the number of steps made. Upon reaching the destination, it will list three asterisks per line. Press `Control+C` to end the attempt.

To use traceroute on a Linux PC

1. Go to a command line prompt.
2. Enter `"/bin/etc/traceroute fortinet.com"`.

The Linux traceroute output is very similar to the MS Windows traceroute output.

Examine routing table contents

The first place to look for information is the routing table.

The routing table is where all the currently used routes are stored for both static and dynamic protocols. If a route is in the routing table, it saves the time and resources of a lookup. If a route isn't used for a while and a new route needs to be added, the oldest least used route is bumped if the routing table is full. This ensures the most recently used routes stay in the table. Note that if your FortiGate unit is in Transparent mode, you are unable to perform this step.

If the FortiGate is running in NAT mode, verify that all desired routes are in the routing table: local subnets, default routes, specific static routes, and dynamic routing protocols.

To check the routing table in the web-based manager, use the Routing Monitor — go to *Router > Monitor > Routing Monitor*. In the CLI, use the command `get router info routing-table all`.

Static routing tips

When your network goes beyond basic static routing, here are some tips to help you plan and manage your static routing.

Always configure a default route

The first thing configured on a router on your network should be the default route. And where possible the default routes should point to either one or very few gateways. This makes it easier to locate and correct problems in the network. By comparison, if one router uses a second router as its gateway which uses a fourth for its gateway and so on, one failure in that chain will appear as an outage for all the devices downstream. By using one or very few addresses as gateways, if there is an outage on the network it will either be very localized or network-wide — either is easy to troubleshoot.

Have an updated network plan

A network plan lists different subnets, user groups, and different servers. Essentially it puts all your resources on the network, and shows how the parts of your network are connected. Keeping your plan updated will also help you troubleshoot problems more quickly when they arise.

A network plan helps your static routing by eliminating potential bottlenecks, and helping troubleshoot any routing problems that come up. Also you can use it to plan for the future and act on any changes to your needs or resources more quickly.

Plan for expansion

No network remains the same size. At some time, all networks grow. If you take future growth into account, there will be less disruption to your existing network when that growth happens. For example allocating a block of addresses for servers can easily prevent having to re-assign IP addresses to multiple servers due to a new server.

With static routing, if you group parts of your network properly you can easily use network masks to address each part of your network separately. This will reduce the amount of administration required both to maintain the routing, and to troubleshoot any problems.

Configure as much security as possible

Securing your network through static routing methods is a good low level method to defend both your important information and your network bandwidth.

- Implement NAT to obscure your IP address is an excellent first step.
- Implement black hole routing to hide which IP addresses are in use or not on your local network.
- Configure and use access control list (ACL) to help ensure you know only valid users are using the network.

All three features limit access to the people who should be using your network, and obscure your network information from the outside world and potential hackers.

Policy routing

Policy routing enables you to redirect traffic away from a static route. This can be useful if you want to route certain types of network traffic differently. You can use incoming traffic's protocol, source address or interface, destination address, or port number to determine where to send the traffic. For example, generally network traffic would go to the router of a subnet, but you might want to direct SMTP or POP3 traffic directly to the mail server on that subnet.

If you have configured the FortiGate unit with routing policies and a packet arrives at the FortiGate unit, the FortiGate unit starts at the top of the Policy Route list and attempts to match the packet with a policy. If a match is found and the policy contains enough information to route the packet (a minimum of the IP address of the next-hop router and the FortiGate interface for forwarding packets to it), the FortiGate unit routes the packet using the information in the policy. If no policy route matches the packet, the FortiGate unit routes the packet using the routing table.



Most policy settings are optional, so a matching policy alone might not provide enough information for forwarding the packet. The FortiGate unit may refer to the routing table in an attempt to match the information in the packet header with a route in the routing table. For example, if the outgoing interface is the only item in the policy, the FortiGate unit looks up the IP address of the next-hop router in the routing table. This situation could happen when the interfaces are dynamic (such as DHCP or PPPoE) and you do not want or are unable to specify the IP address of the next-hop router.

Policy route options define which attributes of an incoming packet cause policy routing to occur. If the attributes of a packet match all the specified conditions, the FortiGate unit routes the packet through the specified interface to the specified gateway.

To view policy routes go to *Router > Static > Policy Routes*.

Create New	Add a policy route. See “Adding a policy route” on page 25 .
Delete icon	Delete the selected policy route.
Edit icon	Edit the selected policy route.
Move To icon	Move the selected policy route. Enter the new position and select OK. For more information, see “Moving a policy route” on page 27 .
#	The ID numbers of configured route policies. These numbers are sequential unless policies have been moved within the table.
Incoming	The interfaces on which packets subjected to route policies are received.
Outgoing	The interfaces through which policy routed packets are routed.
Source	The IP source addresses and network masks that cause policy routing to occur.
Destination	The IP destination addresses and network masks that cause policy routing to occur.

Adding a policy route

To add a policy route, go to *Router > Static > Policy Route* and select *Create New*.

Protocol	<p>Enter the protocol number to match. The Internet Protocol Number is found in the IP packet header. RFC 5237 describes protocol numbers and you can find a list of the assigned protocol numbers here. The range is from 0 to 255. A value of 0 disables the feature.</p> <p>Commonly used <i>Protocol</i> settings include 6 for TCP sessions, 17 for UDP sessions, 1 for ICMP sessions, 47 for GRE sessions, and 92 for multicast sessions.</p>
Incoming Interface	Select the name of the interface through which incoming packets subjected to the policy are received.
Source Address / Mask	To perform policy routing based on IP source address, type the source address and network mask to match. A value of 0.0.0.0/0.0.0.0 disables the feature.
Destination Address / Mask	To perform policy routing based on the IP destination address of the packet, type the destination address and network mask to match. A value of 0.0.0.0/0.0.0.0 disables the feature.
Destination Ports	<p>To perform policy routing based on the port on which the packet is received, type the same port number in the From and To fields. To apply policy routing to a range of ports, type the starting port number in the From field and the ending port number in the To field. A value of 0 disables this feature.</p> <p>The Destination Ports fields are only used for TCP and UDP protocols. The ports are skipped over for all other protocols.</p>
Type of Service	Use a two digit hexadecimal bit pattern to match the service, or use a two digit hexadecimal bit mask to mask out. For more information, see “ Type of Service ” on page 26.
Outgoing Interface	Select the name of the interface through which packets affected by the policy will be routed.
Gateway Address	Type the IP address of the next-hop router that the FortiGate unit can access through the specified interface.

Example policy route

Configure the following policy route to send all FTP traffic received at `port1` out the `port10` interface and to a next hop router at IP address `172.20.120.23`. To route FTP traffic set protocol to 6 (for TCP) and set both of the destination ports to 21, the FTP port.

Protocol	6
Incoming interface	port1
Source address / mask	0.0.0.0/0.0.0.0
Destination address / mask	0.0.0.0/0.0.0.0

Destination Ports	From 21 to 21
Type of Service	bit pattern: 00 (hex) bit mask: 00 (hex)
Outgoing interface	port10
Gateway Address	172.20.120.23

Type of Service

Type of service (TOS) is an 8-bit field in the IP header that enables you to determine how the IP datagram should be delivered, with such qualities as delay, priority, reliability, and minimum cost.

Each quality helps gateways determine the best way to route datagrams. A router maintains a ToS value for each route in its routing table. The lowest priority TOS is 0, the highest is 7 - when bits 3, 4, and 5 are all set to 1. The router tries to match the TOS of the datagram to the TOS on one of the possible routes to the destination. If there is no match, the datagram is sent over a zero TOS route.

Using increased quality may increase the cost of delivery because better performance may consume limited network resources. For more information, see RFC 791 and RFC 1349.

Table 2: The role of each bit in the IP header TOS 8-bit field

bits 0, 1, 2	Precedence	Some networks treat high precedence traffic as more important traffic. Precedence should only be used within a network, and can be used differently in each network. Typically you do not care about these bits.
bit 3	Delay	When set to 1, this bit indicates low delay is a priority. This is useful for such services as VoIP where delays degrade the quality of the sound.
bit 4	Throughput	When set to 1, this bit indicates high throughput is a priority. This is useful for services that require lots of bandwidth such as video conferencing.
bit 5	Reliability	When set to 1, this bit indicates high reliability is a priority. This is useful when a service must always be available such as with DNS servers.
bit 6	Cost	When set to 1, this bit indicates low cost is a priority. Generally there is a higher delivery cost associated with enabling bits 3,4, or 5, and bit 6 indicates to use the lowest cost route.
bit 7	Reserved for future use	Not used at this time.

For example, if you want to assign low delay, and high reliability, say for a VoIP application where delays are unacceptable, you would use a bit pattern of `xxx1x1xx` where an 'x' indicates that bit can be any value. Since all bits are not set, this is a good use for the bit mask; if the mask is set to `0x14`, it will match any TOS packets that are set to low delay and high reliability.

Moving a policy route

A routing policy is added to the bottom of the routing table when it is created. If you prefer to use one policy over another, you may want to move it to a different location in the routing policy table.

The option to use one of two routes happens when both routes are a match, for example 172.20.0.0/255.255.0.0 and 172.20.120.0/255.255.255.0. If both of these routes are in the policy table, both can match a route to 172.20.120.112 but you consider the second one as a better match. In that case the best match route should be positioned before the other route in the policy table.

To change the position of a policy route in the table, go to *Router > Static > Policy Routes* and select *Move To* for the policy route you want to move.

Before/After	Select Before to place the selected Policy Route before the indicated route. Select After to place it following the indicated route.
Policy route ID	Enter the Policy route ID of the route in the Policy route table to move the selected route before or after.

Transparent mode static routing

FortiOS operating modes allow you to change the configuration of your FortiGate unit depending on the role it needs to fill in your network.

NAT/Route operating mode is the standard mode where all interfaces are accessed individually, and traffic can be routed between ports to travel from one network to another.

In transparent operating mode, all physical interfaces act like one interface. The FortiGate unit essentially becomes a bridge — traffic coming in over any interface is broadcast back out over all the interfaces on the FortiGate unit.

In transparent mode, there is no entry for routing at the main level of the menu on the web-based manager display as there is in NAT/Route mode. Routing is instead accessed through the network menu option.

To view the routing table in transparent mode, go to *System > Network > Routing Table*.

When viewing or creating a static route entry in transparent mode there are only three fields available.

Destination IP/Mask	The destination of the traffic being routed. The first entry is attempted first for a match, then the next, and so on until a match is found or the last entry is reached. If no match is found, the traffic will not be routed. Use 0.0.0.0 to match all traffic destinations. This is the default route.
Gateway	Specifies the next hop for the traffic. Generally the gateway is the address of a router on the edge of your network.
Priority	The priority is used if there is more than one match for a route. This allows multiple routes to be used, with one preferred. If the preferred route is unavailable the other routes can be used instead. Valid range of priority can be from 0 to 4 294 967 295. If more than one route matches and they have the same priority it becomes an ECMP situation and traffic is shared among those routes. See “Route priority” on page 20 .

When configuring routing on a FortiGate unit in transparent mode, remember that all interfaces must be connected to the same subnet. That means all traffic will be coming from and leaving on the same subnet. This is important because it limits your static routing options to only the gateways attached to this subnet. For example, if you only have one router connecting your network to the Internet then all static routing on the FortiGate unit will use that gateway. For this reason static routing on FortiGate units in transparent mode may be a bit different, but it is not as complex as routing in NAT/Route mode.

Static routing example

This is an example of a typical small network configuration that uses only static routing.

This network is in a dentist office that includes a number of dentists, assistants, and office staff. The size of the office is not expected to grow significantly in the near future, and the network usage is very stable—there are no new applications being added to the network.

The users on the network are:

- admin staff - access to local patient records, and go online for billing purposes
- dentists - access and update local patient records, research online from desk
- assistants - access and update local patient records in exam rooms

The distinction here is mainly that only the admin staff and dentist's office proper need access to the internet—all the other traffic is local and doesn't need to need to leave the local network. Routing is only required for the outbound traffic, and the computers that have valid outbound traffic.



Only configuring routing on computers that need it will act as an additional layer of security by helping prevent malicious traffic from leaving the network.

This section includes the following topics:

- [Network layout and assumptions](#)
- [General configuration steps](#)
- [Configure FortiGate unit](#)
- [Configure Admin PC and Dentist PCs](#)
- [Testing network configuration](#)

Network layout and assumptions

The computers on the network are admin staff computers, dentist office computers, and dental exam room computers. While there are other devices on the local network such as printers, they do not need Internet access or any routing.

This networked office equipment includes 1 admin staff PC, 3 dentist PCs, and 5 exam room PCs. There are also a network printer, and a router on the network as well.

Assumptions about these computers, and network include:

- the FortiGate unit is a model with interfaces labeled port1 and port2
- the FortiGate unit has been installed and is configured in NAT/Route mode
- VDOMs are not enabled
- the computers on the network are running MS Windows software
- any hubs required in the network are not shown in the network diagram
- the network administrator has access to the ISP IP addresses, and is the super_admin administrator on the FortiGate unit




Table 3: Static routing example device names, IP addresses, and level of access

Device Name(s)	IP address	Need external access?
Router	192.168.10.1	YES
Admin	192.168.10.11	YES
Dentist1-3	192.168.10.21-23	YES
Exam1-5	192.168.10.31-35	NO
Printer	192.168.10.41	NO

General configuration steps

The steps to configuring routing on this network are:

1. Get your ISP information such as DNS, gateway, etc.
2. Configure FortiGate unit
3. Configure Admin PC and Dentist PCs
4. Testing network configuration

Get your ISP information such as DNS, gateway, etc.

Your local network connects to the Internet through your Internet Service Provider (ISP). They have IP addresses that you need to configure your network and routing.

The addresses needed for routing are your assigned IP address, DNS servers, and the gateway.

Configure FortiGate unit

The FortiGate unit will have two interfaces in use—one connected to the internal network and one connected to the external network. Port1 will be the internal interface, and port2 will be the external interface.

To configure the FortiGate unit:

1. Configure the internal interface (port1)
2. Configure the external interface (port2)
3. Configure networking information
4. Configure basic security policies
5. Configure static routing

Configure the internal interface (port1)

To configure the internal interface (port1) - web based manager

1. Go to *System > Network > port1* and select *Edit*.
2. Enter the following:

Addressing Mode	Manual
IP/Netmask	172.100.1.1/255.255.255.0
Administrative Access	HTTPS, PING, TELNET
Description	Internal network

To configure the internal interface (port1) - CLI

```
config system interface
  edit port1
    set IP 192.168.10.1 255.255.255.0
    set allowaccess https ping telnet
    set description "internal network"
  end
end
```

Configure the external interface (port2)

The external interface connects to your ISP's network. You need to know the IP addresses in their network that you should connect to. Use their addresses when you get them, however for this example we will assume the address your ISP gave you is 172.100.20.20 will connect to the gateway at 172.100.20.5 on their network, and their DNS servers are 172.11.22.33 and 172.11.22.34.

To configure the internal interface (port2) - web based manager

1. Go to *System > Network > port2* and select *Edit*.

2. Enter the following:

Addressing Mode	Manual
IP/Netmask	172.100.20.20/255.255.255.0
Administrative Access	HTTPS, PING, TELNET
Description	Internal network

To configure the internal interface (port2) - CLI

```
configure system interface
  edit port2
    set IP 172.100.20.20 255.255.255.0
    set allowaccess https ping telnet
    set description "internal network"
  end
end
```

Configure networking information

Networking information includes the gateway, and DNS servers. Your FortiGate unit requires a connection to the Internet for antivirus and other periodic updates.

To configure networking information - web-based manager

1. Go to *System > Network > Options*.
2. Enter the primary and secondary DNS addresses.
3. Select *Apply*.

To configure networking information - CLI

```
config system global
  set dns_1 172.11.22.33
  set dns_2 172.11.22.34
end
```

Configure basic security policies

For traffic to flow between the internal and external ports in both directions, two security policies are required as a minimum. More can be used to farther limit or direct traffic as needed, but will not be included here.

Before configuring the security policies, a firewall address group is configured for the PCs that are allowed Internet access. This prevents PC without Internet privileges from accessing the Internet.

The security policy assumptions are:

- Only the basic networking services have been listed as allowed for added security. Others can easily be added as the users require them.
- In this example to keep things simple, both incoming and outgoing security policies are the same. In a real network there are applications that are allowed out but not in, and vice versa.
- Endpoint control has been enabled to ensure that all computers on the local network are running FortiClient and those installs are up to date. This feature ensures added security on your local network without the need for the network administrator to continually bother users

to update their software. The FortiGate unit can store an up to date copy of the FortiClient software and offer a URL to it for users to install it if they need to.

To configure security policies - web-based manager

1. Go to *Firewall Objects > Address*.
2. Create a new Firewall Address entry for each of:

PC Name	IP Address	Interface
Admin	192.168.10.11	port1
Dentist1	192.168.10.21	port1
Dentist2	192.168.10.22	port1
Dentist3	192.168.10.23	port1

3. Go to *Firewall Objects > Address > Groups*.
4. Select *Create New*.
5. Name the group *Internet_PCs*.
6. Select *Admin, Dentist1, Dentist2, and Dentist3* to be members of the group.
7. Select *OK*.
8. Go to *Policy > Policy > Policy*.
9. Select *Create New*.
10. Enter the following: *DH - port2(external) -> port1(internal)*

Source Interface/Zone	port2
Source Address	all
Destination Interface/Zone	port1
Destination Address	Internet_PCs
Schedule	always
Service	Multiple. Select <i>DHCP, DNS, FTP, HTTP, HTTPS, NTP, POP3, SMTP, SSH</i> .
Action	ACCEPT
Enable Endpoint Control Check	Enabled Enforce FortiClient AV Up-to-date
Log Allowed Traffic	Enabled

11. Select *OK*.
12. Select *Create New*.
13. Enter the following:

Source Interface/Zone	port1
Source Address	Internet_PCs
Destination Interface/Zone	port2
Destination Address	all
Schedule	always
Service	Multiple. Select <i>DHCP, DNS, FTP, HTTP, HTTPS, NTP, POP3, SMTP, SSH.</i>
Action	ACCEPT
Enable Endpoint Control Check	Enabled Enforce FortiClient AV Up-to-date
Log Allowed Traffic	Enabled

14. Select OK.

To configure security policies - CLI

```

config firewall address
  edit "Admin"
    set associated-interface "port1"
    set subnet 192.168.10.11 255.255.255.255
  next
  edit "Dentist1"
    set associated-interface "port1"
    set subnet 192.168.10.21 255.255.255.255
  next
  edit "Dentist2"
    set associated-interface "port1"
    set subnet 192.168.10.22 255.255.255.255
  next
  edit "Dentist3"
    set associated-interface "port1"
    set subnet 192.168.10.23 255.255.255.255
  end
config firewall addrgrp
  edit Internet_PCs
    set member Admin Dentist1 Dentist2 Dentist3
  end
config firewall policy
  edit 1
    set srcintf port1
    set dstintf port2
    set srcaddr Internet_PCs

```

```

set dstaddr all
set action accept
set schedule always
set service "DHCP" "DNS" "FTP" "HTTP" "HTTPS" "NTP" "POP3"
  "SMTP" "SSH"
set logtraffic enable
set endpoint-check enable
set label "Section2"
set endpoint-restrict-check no-av db-outdated
next
edit 2
  set srcintf port2
  set dstintf port1
  set srcaddr all
  set dstaddr Internet_PCs
  set action accept
  set schedule always
  set service "DHCP" "DNS" "FTP" "HTTP" "HTTPS" "NTP" "POP3"
    "SMTP" "SSH"
  set logtraffic enable
  set endpoint-check enable
  set label "Section2"
  set endpoint-restrict-check no-av db-outdated
end
end

```

Configure static routing

With the rest of the FortiGate unit configured, static routing is the last step before moving on to the rest of the local network. All traffic on the local network will be routed according to this static routing entry.

To configure Fortinet unit static routing - web-based manager

1. Go to *Routing > Static*.
2. Select *Edit* for the top route on the page.
3. Enter the following:

Destination IP/Mask	172.100.20.5
Device	port2
Gateway	172.100.20.5
Distance	10

4. Select *OK*.

To configure Fortinet unit static routing - CLI

```
configure routing static
  edit 1
    set gateway 172.100.20.5
    set distance 10
    set device port2
    set dst 0.0.0.0
  end
end
```

Configure Admin PC and Dentist PCs

With the router configured, next we need to configure the computers that need Internet access. These computers need routing to be configured on them. As the other computers do not require routing, they are not included here.

The procedure to configure these computers is the same. Repeat the following procedure for the corresponding PCs.



The Windows CLI procedure does not configure the DNS entries. It just adds the static routes.

To configure routing and DNS on Admin and Dentist PCs - Windows GUI

1. On PC, select *Start > Control Panel > Network Connections*.
2. Right click on the network connection to your local network that has a status of *Connected*, and select *Properties*.
3. Under the *General* tab, from the list select *TCP/IP*, and *Properties*.
4. Under *Gateway*, enter the FortiGate unit address (192.168.10.1).
5. Enter the primary and secondary DNS server addresses from your ISP (172.11.22.33 and 172.11.22.34).
6. Select *OK*.

To configure routing on Admin and Dentist PCs - Windows CLI

1. On PC, select *Start > Run*, enter "cmd", and select *OK*.
2. At the command prompt, type

```
route ADD 0.0.0.0 MASK 0.0.0.0 172.100.20.5 METRIC 10
route ADD 192.168.10.0 MASK 255.255.255.0 192.168.10.1 METRIC 5
```
3. Confirm these routes have been added. Type:

```
route PRINT
```

If you do not see the two routes you added, try adding them again paying attention to avoid spelling mistakes.
4. Test that you can communicate with other computers on the local network, and with the Internet. If there are no other computers on the local network, connect to the FortiGate unit.

Configure other PCs on the local network

The PCs on the local network without Internet access (the exam room PCs) can be configured now.

As this step does not require any routing, details have not been included.

Testing network configuration

There are three tests to run on the network to ensure proper connectivity.

- To test that PCs on the local network can communicate
- Test that Internet_PCs on the local network can access the Internet.
- Test that non-Internet_PCs can not access the Internet.

To test that PCs on the local network can communicate

1. Select any two PCs on the local network, such as Exam4 and Dentist3.
2. On the Exam4 PC, at the command prompt enter `ping 192.168.10.23`.

The output from this command should appear similar to the following.

```
Pinging 192.168.10.23 with 32 bytes of data:
```

```
Reply from 192.168.10.23: bytes=32 time<1m TTL=255
Reply from 192.168.10.23: bytes=32 time<1m TTL=255
Reply from 192.168.10.23: bytes=32 time<1m TTL=255
```

3. At the command prompt enter `exit` to close the window.
4. On the Dentist3 PC, at the command prompt enter `ping 192.168.10.34`.

The output from this command should appear similar to the following.

```
Pinging 192.168.10.34 with 32 bytes of data:
```

```
Reply from 192.168.10.34: bytes=32 time<1m TTL=255
Reply from 192.168.10.34: bytes=32 time<1m TTL=255
Reply from 192.168.10.34: bytes=32 time<1m TTL=255
```

5. At the command prompt enter `exit` to close the window.
6. Repeat these steps for all PCs on the local network.

If the output does not appear similar to above, there is a problem with the network configuration between these two PCs.

To test that Internet_PCs can reach the Internet

The easiest way to access the internet is with an Internet browser. However, if that doesn't work its best to do a traceroute to see at what point the problem is. This can help determine if it is a networking problem such as cabling, or if its an access problem such as this PC not having Internet access.

1. Select any PC on the local network that is supposed to have Internet access, such as Admin.
2. On the Admin PC, open an Internet browser and attempt to access a website on the Internet such as <http://www.fortinet.com>.

If this is successful, this PC has Internet access.

3. If step 2 was not successful, at the command prompt on the PC enter `tracert 22.11.22.33`.

The output from this command should appear similar to:

Pinging 22.11.22.33 with 32 bytes of data:

```
Reply from 22.11.22.33: bytes=32 time<1m TTL=255
```

```
Reply from 22.11.22.33: bytes=32 time<1m TTL=255
```

```
Reply from 22.11.22.33: bytes=32 time<1m TTL=255
```

Advanced static example: ECMP failover and load balancing

Equal Cost Multi-Path (ECMP) load balancing and failover are methods that extend basic static routing. They allow you to use your network bandwidth more effectively and with less down time than if you used basic static routing alone.

The concepts in this section include:

- Equal-Cost Multi-Path (ECMP)
- Configuring interface status detection for gateway load balancing
- Configuring spillover or usage-based ECMP
- Configuring weighted static route load balancing

Equal-Cost Multi-Path (ECMP)

FortiOS uses equal-cost multi-path (ECMP) to distribute traffic to the same destination such as the Internet or another network. Using ECMP you can add multiple routes to the destination and give each of those routes the same distance and priority.



If multiple routes to the same destination have the same priority but different distances, the route with the lowest distance is used. If multiple routes to the same destination have the same distance but different priorities, the route with the lowest priority is used. Distance takes precedence over priority. If multiple routes to the same destination have different distances and different priorities, the route with the lowest distance is always used even if it has the highest priority.

If more than one ECMP route is available, you can configure how the FortiGate unit selects the route to be used for a communication session. If only one ECMP route is available (for example, because an interface cannot process traffic because interface status detection does not receive a reply from the configured server) then all traffic uses this route.

Previous versions of FortiOS provided source IP-based load balancing for ECMP routes, but now FortiOS includes three configuration options for ECMP route failover and load balancing:

**Source based
(also called source IP
based)**

The FortiGate unit load balances sessions among ECMP routes based on the source IP address of the sessions to be load balanced. This is the default load balancing method. No configuration changes are required to support source IP load balancing.

Weighted Load Balance (also called weight-based)

The FortiGate unit load balances sessions among ECMP routes based on weights added to ECMP routes. More traffic is directed to routes with higher weights. After selecting weight-based you must add weights to static routes.

Spillover (also called usage-based)

The FortiGate unit distributes sessions among ECMP routes based on how busy the FortiGate interfaces added to the routes are.

After selecting spill-over you add route Spillover Thresholds to interfaces added to ECMP routes. The FortiGate unit sends all ECMP-routed sessions to the lowest numbered interface until the bandwidth being processed by this interface reaches its spillover threshold. The FortiGate unit then spills additional sessions over to the next lowest numbered interface.

The Spillover Thresholds range is 0-2097000 KBps.

You can configure only one of these ECMP route failover and load balancing methods in a single VDOM. If your FortiGate unit is configured for multiple VDOM operation, each VDOM can have its own ECMP route failover and load balancing configuration.

To configure the ECMP load balancing method from the web-based manager

1. Go to *Router > Static > Settings*.
2. Set *ECMP Load Balance Method* to *Source IP based*, *Weighted Load Balance*, or *Spillover*.

To configure the ECMP load balancing method from the CLI

For example, to set the load balancing method to usage-based, enter the following:

```
config system settings
    set v4-ecmp-mode usage-based
end
```

ECMP routing of simultaneous sessions to the same destination IP address

When the FortiGate unit selects an ECMP route for a session, a route cache is created that matches the route with the destination IP address of the session. All new sessions to the same destination IP address use the same route until the route is flushed from the cache. Routes are flushed from the cache after a period of time when no new sessions to the destination IP address are received.


The route cache improves FortiGate unit routing performance by reducing how often the FortiGate unit looks up routes in the routing table.

If the FortiGate unit receives a large number of sessions with the same destination IP address, because all of these sessions will be processed by the same route, it may appear that sessions are not distributed according to the ECMP route failover and load balancing configuration.

Configuring interface status detection for gateway load balancing

Interface status detection is used for ECMP route failover and load balancing. Interface status detection consists of the unit confirming that packets sent from an interface result in a response from a server. You can use up to three different protocols to confirm that an interface can connect to the server. Usually the server is the next-hop router that leads to an external network or the Internet. Interface status detection sends a packet using the configured protocols. If a response is received from the server, the unit assumes the interface can connect to the network. If a response is not received, the unit assumes that the interface cannot connect to the network.

Since it is possible that a response may not be received, even if the server and the network are operating normally, the dead gateway detection configuration controls the time interval between testing the connection to the server and the number of times the test can fail before the unit assumes that the interface cannot connect to the server.

	<p>As long as the unit receives responses for at least one of the protocols that you select, the unit assumes the server is operating and can forward packets. Responding to more than one protocol does not enhance the status of the server or interface.</p>
---	---

To configure gateway failover detection for an interface

1. Go to *Router > Static > Settings*.
2. Under *Dead Gateway Detection*, select *Create New*.
3. Enter the following information:

Interface	Select the interface to test.
Gateway IP	Enter the IP address of the gateway.
Ping Server	Enter the IP address of the server to test.
Detect Protocol	Select one of the following protocols.
ICMP Ping	Use standard ICMP ping to confirm that the server is responding. Ping confirms that the server can respond to an ICMP ping request.
TCP Echo	<p>Use TCP echo to confirm that the server is responding. Select this option if the server is configured to provide TCP echo services. In some cases a server may be configured to reply to TCP echo requests but not to reply to ICMP pings.</p> <p>TCP echo uses TCP packets on port number 7 to send a text string to the server and expect an echo reply back from the server. The echo reply just echoes back the same text to confirm that the server can respond to TCP requests.</p> <p>FortiGate units do not recognize RST (reset) packets from TCP Echo servers as normal TCP echo replies. If the unit receives an RST response to a TCP echo request, the unit assumes the server is unreachable.</p>
UDP Echo	<p>Use UDP echo to detect the server. Select this option if the server is configured to provide UDP echo services. In some cases a server may be configured to reply to UDP echo requests but not to reply ICMP pings.</p> <p>UDP echo uses UDP packets on port number 7 to send a text string to the server and expects an echo reply from the server. The echo reply just echoes back the same text to confirm that the server can respond to UDP requests.</p>
Ping Interval	Enter the interval between pings, in seconds.
Failover Threshold	Enter the number of times the test can fail before the unit assumes that the interface cannot connect to the server.
HA Priority	Set the HA priority, if configuring an HA cluster.

4. Select *OK*.

To configure gateway failover detection for an interface - CLI

```
config router gwdetect
  edit port1
    set protocol ping
    set server 10.10.10.1
    set interval 5
    set failtime 5
  end
```

Configuring spillover or usage-based ECMP

Spill-over or usage-based ECMP routes new sessions to interfaces that have not reached a configured bandwidth limit (called the *Spillover Threshold* or a route-spillover threshold). To configure spill-over or usage-based ECMP routing, you enable spill-over ECMP, add ECMP routes, and add a *Spillover Threshold* to the interfaces used by the ECMP routes. Set the *Spillover Thresholds* to limit the amount of bandwidth processed by each interface. The range is 0 to 2 097 000 Kbps. The threshold counts only outgoing traffic.

With spill-over ECMP routing configured, the FortiGate unit routes new sessions to an interface used by an ECMP route until that interface reaches its *Spillover Threshold*. Then, when the threshold of that interface is reached, new sessions are routed to one of the other interfaces used by the ECMP routes.

To add Spillover Thresholds to interfaces - web-based manager

Use the following steps to enable usage based ECMP routing, add Spillover Thresholds to FortiGate interfaces `port3` and `port4`, and then to configure EMCP routes with device set to `port3` and `port4`.

1. Go to *Router > Static > Settings*.
2. Set *ECMP Load Balance Method* to *Spillover*.
3. Go to *Router > Static > Static Routes*.
4. Add ECMP routes for `port3` and `port4`.

Destination IP/Mask	192.168.20.0/24
Device	port3
Gateway	172.20.130.3
Advanced	
Distance	10

Destination IP/Mask	192.168.20.0/24
Device	port4
Gateway	172.20.140.4
Advanced	
Distance	10

5. Go to *System > Network > Settings*.
6. Edit `port3` and `port4` and add the following spillover-thresholds:

Interface	port3
Spillover Threshold	100

Interface	port4
Spillover Threshold	200

To add Spillover Thresholds to interfaces - CLI

```

config system settings
    set v4-ecmp-mode usage-based
end
config router static
    edit 1
        set device port3
        set dst 192.168.20.0 255.255.255.0
        set gateway 172.20.130.3
    next
    edit 2
        set device port4
        set dst 192.168.20.0 255.255.255.0
        set gateway 172.20.140.4
    end
config system interface
    edit port3
        set spillover-threshold 100
    next
    edit port4
        set spillover-threshold 200
    end

```

Detailed description of how spill-over ECMP selects routes

When you add ECMP routes they are added to the routing table in the order displayed by the routing monitor or by the `get router info routing-table static` command. This order is independent of the configured bandwidth limit.

The FortiGate unit selects an ECMP route for a new session by finding the first route in the routing table that sends the session out a FortiGate unit interface that is not processing more traffic than its configured route spill-over limit.



A new session to a destination IP address that already has an entry in the routing cache is routed using the route already added to the cache for that destination address. See [“ECMP routing of simultaneous sessions to the same destination IP address”](#) on page 38.

For example, consider a FortiGate unit with interfaces port3 and port4 both connected to the Internet through different ISPs. ECMP routing is set to usage-based and route spillover for to 100 Kbps for port3 and 200 Kbps for port4. Two ECMP default routes are added, one for port3 and one for port4.

If the route to port3 is higher in the routing table than the route to port4, the FortiGate unit sends all default route sessions out port3 until port3 is processing 100Kbps of data. When port3 reaches its configured bandwidth limit, the FortiGate unit sends all default route sessions out

port4. When the bandwidth usage of port3 falls below 100Kbps, the FortiGate again sends all default route sessions out port3.

New sessions with destination IP addresses that are already in the routing cache; however, use the cached routes. This means that even if port3 is exceeding its bandwidth limit, new sessions can continue to be sent out port3 if their destination addresses are already in the routing cache. As a result, new sessions are sent out port4 only if port3 exceeds its bandwidth limit and if the routing cache does not contain a route for the destination IP address of the new session.

Also, the switch over to port4 does not occur as soon as port3 exceeds its bandwidth limit. Bandwidth usage has to exceed the limit for a period of time before the switch over takes place. If port3 bandwidth usage drops below the bandwidth limit during this time period, sessions are not switched over to port4. This delay reduces route flapping.

FortiGate usage-based ECMP routing is not actually load balancing, since routes are not distributed evenly among FortiGate interfaces. Depending on traffic volumes, most traffic would usually be processed by the first interface with only spillover traffic being processed by other interfaces.

If you are configuring usage-based ECMP in most cases you should add spillover thresholds to all of the interfaces with ECMP routes. The default spillover threshold is 0 which means no bandwidth limiting. If any interface has a spillover threshold of 0, no sessions will be routed to interfaces lower in the list unless the interface goes down or is disconnected. An interface can go down if *Detect interface status for Gateway Load Balancing* does not receive a response from the configured server.

Determining if an interface has exceeded its Spillover Threshold

You can use the `diagnose netlink dstmac list` CLI command to determine if an interface is exceeding its Spillover Threshold. If the command displays `over_bps=1` the interface is exceeding its threshold. If `over_bps=0` the interface has not exceeded its threshold.

```
dev=Wifi mac=00:00:00:00:00:00 src-vis-os src-vis-host src-vis-user
rx_tcp_mss=0 tx_tcp_mss=0 overspill-threshold=0 bytes=0 over_bps=0
sampler_rate=0
```

Configuring weighted static route load balancing

Configure weighted load balancing to control how the FortiGate unit distributes sessions among ECMP routes by adding weights for each route. Add higher weights to routes that you want to load balance more sessions to.

With the ECMP load balancing method set to weighted, the FortiGate unit distributes sessions with different destination IPs by generating a random value to determine the route to select. The probability of selecting one route over another is based on the weight value of each route. Routes with higher weights are more likely to be selected.

Large numbers of sessions are evenly distributed among ECMP routes according to the route weight values. If all weights are the same, sessions are distributed evenly. The distribution of a small number of sessions; however, may not be even. For example, its possible that if there are two ECMP routes with the same weight; two sessions to different IP addresses could use the same route. On the other hand, 10,000 sessions with different destination IPs should be load balanced evenly between two routes with equal rates. The distribution could be 5000:5000 or 50001:4999. Also, 10 000 sessions with different destination IP addresses should be load balanced as 3333:6667 if the weights for the two routes are 100 and 200.

Weights only affect how routes are selected for sessions to new destination IP addresses. New sessions to IP addresses already in the routing cache are routed using the route for the session already in the cache. So in practice sessions will not always be distributed according to the routing weight distribution.

To add weights to static routes from the web-based manager

1. Go to *Router > Static > Settings*.
2. Set *ECMP Load Balance Method* to *Weighted Load Balance*.
3. Go to *Router > Static > Static Routes*.
4. If needed, add new static routes, for example:

Destination IP/Mask	192.168.20.0/24
Device	port1
Gateway	172.20.110.1
Distance	10

Destination IP/Mask	192.168.20.0/24
Device	port2
Gateway	172.20.120.2
Distance	10

5. Go to *Router > Static > Settings*.
6. Select a number next to an interface name, and choose *Edit* to change it.

For example, set the weight of port1 to 100 and the weight of port2 to 200.

Dynamic Routing Overview

This section provides an overview of dynamic routing, and how it compares to static routing. For details on various dynamic routing protocols, see the following chapters for detailed information.

The following topics are included in this section:

- [What is dynamic routing?](#)
- [Comparison of dynamic routing protocols](#)
- [Choosing a routing protocol](#)
- [Dynamic routing terminology](#)
- [IPv6 in dynamic routing](#)

What is dynamic routing?

Dynamic routing uses a dynamic routing protocol to automatically select the best route to put into the routing table. So instead of manually entering static routes in the routing table, dynamic routing automatically receives routing updates, and dynamically decides which routes are best to go into the routing table. It's this intelligent and hands-off approach that makes dynamic routing so useful.

Dynamic routing protocols vary in many ways and this is reflected in the various administrative distances assigned to routes learned from dynamic routing. These variations take into account differences in reliability, speed of convergence, and other similar factors. For more information on these administrative distances, see [“Multipath routing and determining the best route”](#) on page 19.

This section includes:

- [Comparing static and dynamic routing](#)
- [Dynamic routing protocols](#)
- [Minimum configuration for dynamic routing](#)

Comparing static and dynamic routing

A common term used to describe dynamic routing is convergence. Convergence is the ability to work around network problems and outages — for the routing to come together despite obstacles. For example, if the main router between two end points goes down, convergence is the ability to find a way around that failed router and reach the destination. Static routing has zero convergence beyond trying the next route in its limited local routing table — if a network administrator doesn't fix a routing problem manually, it may never be fixed, resulting in a downed network. Dynamic routing solves this problem by involving routers along the route in the decision-making about the optimal route, and using the routing tables of these routers for potential routes around the outage. In general, dynamic routing has better scalability, robustness, and convergence. However, the cost of these added benefits include more

complexity and some overhead: the routing protocol uses some bandwidth for its own administration.

Table 4: Comparing static and dynamic routing

Feature	Static Routing	Dynamic Routing
Hardware support	Supported by all routing hardware	May require special, more expensive routers
Router Memory Required	Minimal	Can require considerable memory for larger tables
Complexity	Simple	Complex
Overhead	None	Varying amounts of bandwidth used for routing protocol updates
Scalability	Limited to small networks	Very scalable, better for larger networks
Robustness	None - if a route fails it has to be fixed manually	Robust - traffic routed around failures automatically
Convergence	None	Varies from good to excellent

Dynamic routing protocols

A dynamic routing protocol is an agreed-on method of routing that the sender, receiver, and all routers along the path (route) support. Typically the routing protocol involves a process running on all computers and routers along that route to enable each router to handle routes in the same way as the others. The routing protocol determines how the routing tables are populated along that route, how the data is formatted for transmission, and what information about a route is included with that route. For example RIP, and BGP use distance vector algorithms, where OSPF uses a shortest path first algorithm. Each routing protocol has different strengths and weaknesses — one protocol may have fast convergence, while another may be very reliable, and a third is very popular for certain businesses like Internet Service Providers (ISPs).

Dynamic routing protocols are different from each other in a number of ways, such as:

- [Classful versus classless routing protocols](#)
- [Interior versus exterior routing protocols](#)
- [Distance vector versus link-state protocols](#)

Classful versus classless routing protocols

Classful or classless routing refers to how the routing protocol handles the IP addresses. In classful addresses there is the specific address, and the host address of the server that address is connected to. Classless addresses use a combination of IP address and netmask.

Classless Inter-Domain Routing (CIDR) was introduced in 1993 (originally with RFC 1519 and most recently with RFC 4632) to keep routing tables from getting too large. With Classful routing, each IP address requires its own entry in the routing table. With Classless routing, a series of addresses can be combined into one entry potentially saving vast amounts of space in routing tables.

Current routing protocols that support classless routing out of necessity include RIPv2, BGP, IS-IS, and OSPF. Older protocols such as RIPv1 do not support CIDR addresses.

Interior versus exterior routing protocols

The names *interior* and *exterior* are very descriptive. Interior routing protocols are designed for use within a contained network of limited size, whereas exterior routing protocols are designed to link multiple networks together. They can be used in combination in order to simplify network administration. For example, a network can be built with only border routers of a network running the exterior routing protocol, while all the routers on the network run the interior protocol, which prevents them from connecting outside the network without passing through the border. Exterior routers in such a configuration must have both exterior and interior protocols, to communicate with the interior routers and outside the network.

Nearly all routing protocols are interior routing protocols. Only BGP is commonly used as an exterior routing protocol.

You may see interior gateway protocol (IGP) used to refer to interior routing protocols, and exterior gateway protocol (EGP) used to refer to exterior routing protocols.

Distance vector versus link-state protocols

Every routing protocol determines the best route between two addresses using a different method. However, there are two main algorithms for determining the best route — Distance vector and Link-state.

Distance vector protocols

In distance vector protocols, routers are told about remote networks through neighboring routers. The distance part refers to the number of hops to the destination, and in more advanced routing protocols these hops can be weighted by factors such as available bandwidth and delay. The vector part determines which router is the next step along the path for this route. This information is passed along from neighboring routers with routing update packets that keep the routing tables up to date. Using this method, an outage along a route is reported back along to the start of that route, ideally before the outage is encountered.

On distance vector protocols, RFC 1058 which defines RIP v1 states the following:

Distance vector algorithms are based on the exchange of only a small amount of information. Each entity (gateway or host) that participates in the routing protocol is assumed to keep information about all of the destinations within the system. Generally, information about all entities connected to one network is summarized by a single entry, which describes the route to all destinations on that network.

There are four main weaknesses inherent in the distance vector method. Firstly, the routing information is not discovered by the router itself, but is instead reported information that must be relied on to be accurate and up-to-date. The second weakness is that it can take a while for the information to make its way to all the routers who need the information — in other words it can have slow convergence. The third weakness is the amount of overhead involved in passing these updates all the time. The number of updates between routers in a larger network can significantly reduce the available bandwidth. The fourth weakness is that distance vector protocols can end up with routing-loops. Routing loops are when packets are routed for ever around a network, and often occur with slow convergence. The bandwidth required by these infinite loops will slow your network to a halt. There are methods of preventing these loops however, so this weakness is not as serious as it may first appear.

Link-state protocols

Link-state protocols are also known as shortest path first protocols. Where distance vector uses information passed along that may or may not be current and accurate, in link-state protocols each router passes along only information about networks and devices directly connected to it. This results in a more accurate picture of the network topology around your router, allowing it to make better routing decisions. This information is passed between routers using link-state advertisements (LSAs). To reduce the overhead, LSAs are only sent out when information

changes, compared to distance vector sending updates at regular intervals even if no information has changed. The more accurate network picture in link-state protocols greatly speed up convergence and avoid problems such as routing-loops.

Minimum configuration for dynamic routing

Dynamic routing protocols do not pay attention to routing updates from other sources, unless you specifically configure them to do so using CLI redistribute commands within each routing protocol.

The minimum configuration for any dynamic routing to function is to have dynamic routing configured on one interface on the FortiGate unit, and one other router configured as well. Some protocols require larger networks to function as designed.

Table 5: Minimum configuration based on dynamic protocol

	BGP	RIP	OSPF / IS-IS
Interface	yes	yes	yes
Network	yes	yes	yes
AS	local and neighbor	no	yes
Neighbors	at least one	at least one	at least one
Version	no	yes	no
Router ID	no	no	yes

Comparison of dynamic routing protocols

Each dynamic routing protocol was designed to meet a specific routing need. Each protocol does some things well, and other things not so well. For this reason, choosing the right dynamic routing protocol for your situation is not an easy task.

Features of dynamic routing protocols

Each protocol is better suited for some situations over others.

Choosing the best dynamic routing protocol depends on the size of your network, speed of convergence required, the level of network maintenance resources available, what protocols the networks you connect to are using, and so on. For more information on these dynamic routing protocols, see [“Routing Information Protocol \(RIP\)” on page 59](#), [“Border Gateway Protocol \(BGP\)” on page 98](#), [“Open Shortest Path First \(OSPF\)” on page 136](#), and [“Intermediate System to Intermediate System Protocol \(IS-IS\)” on page 178](#).

Table 6: Comparing RIP, BGP, and OSPF dynamic routing protocols

Protocol	RIP	BGP	OSPF / IS-IS
Routing algorithm	Distance Vector, basic	Distance Vector, advanced	Link-state
Common uses	Small non-complex networks	Network backbone, ties multinational offices together	Common in large, complex enterprise networks

Table 6: Comparing RIP, BGP, and OSPF dynamic routing protocols (Continued)

Protocol	RIP	BGP	OSPF / IS-IS
Strengths	Fast and simple to implement Near universal support Good when no redundant paths	Graceful restart BFD support Only needed on border routers Summarize routes	Fast convergence Robust Little management overhead No hop count limitation Scalable
Weakness	Frequent updates can flood network Slow convergence Maximum 15 hops may limit network configuration	Required full mesh in large networks can cause floods Route flap Load-balance multi-homed networks Not available on low-end routers	Complex No support for unequal cost multipath routing Route summary can require network changes
Authentication	Optional authentication using text string or MD5 password. (RIP v1 has no authentication)		
IPv6 Support	Only in RIPng	Only in BGP4+	Only in OSPF6 / Integrated IS-IS

Routing protocols

Routing Information Protocol (RIP) uses classful routing, as well as incorporating various methods to stop incorrect route information from propagating, such as the poisoned horizon method. However, on larger networks its frequent updates can flood the network and its slow convergence can be a problem.

Border Gateway Protocol (BGP) has been the core Internet backbone routing protocol since the mid 1990s, and is the most used interior gateway protocol (IGP). However, some configurations require full mesh connections which flood the network, and there can be route flap and load balancing issues for multihomed networks.

Open Shortest Path First (OSPF) is commonly used in large enterprise networks. It is the protocol of choice mainly due to its fast convergence. However, it can be complicated to setup properly.

Intermediate System to Intermediate System (IS-IS) Protocol allows routing of ISO's OSI protocol stack Connectionless Network Service (CLNS). IS-IS is an Interior Gateway Protocol (IGP) not intended to be used between Autonomous Systems (ASes). IS-IS is a link state protocol well-suited to smaller networks that is in widespread use and has near universal support on routing hardware.

Multicast addressing is used to broadcast from one source to many destinations efficiently. Protocol Independent Multicast (PIM) is the protocol commonly used in enterprises, multimedia content delivery, and stock exchanges.

Routing algorithm

Each protocol uses a slightly different algorithm for choosing the best route between two addresses on the network. The algorithm is the “intelligent” part of a dynamic protocol because the algorithm is responsible for deciding which route is best and should be added to the local routing table. RIP and BGP use distance vector algorithms, where OSPF and IS-IS use link-state or a shortest path first algorithm.

Vector algorithms are essentially based on the number of hops between the originator and the destination in a route, possibly weighting hops based on how reliable, fast, and error-free they are.

The link-state algorithm used by OSPF and IS-IS is called the Dijkstra algorithm. Link-state treats each interface as a link, and records information about the state of the interface. The Dijkstra algorithm creates trees to find the shortest paths to the routes it needs based on the total cost of the parts of the routes in the tree.

For more information on the routing algorithm used, see [“Distance vector versus link-state protocols”](#) on page 46.

Authentication

If an attacker gains access to your network, they can masquerade as a router on your network to either gain information about your network or disrupt network traffic. If you have a high quality firewall configured, it will help your network security and stop many of this type of threat. However, the main method for protecting your routing information is to use authentication in your routing protocol. Using authentication on your FortiGate unit and other routers prevents access by attackers — all routers must authenticate with passwords, such as MD5 hash passwords, to ensure they are legitimate routers.

When configuring authentication on your network, ensure you configure it the same on all devices on the network. Failure to do so will create errors and outages as those forgotten devices fail to connect to the rest of the network.

For example, to configure an MD5 key of 123 on an OSPF interface called `ospf_test`, enter the following CLI command:

```
config router ospf
  config ospf-interface
    edit ospf_test
      set authentication md5
      set md5-key 123
    end
  end
end
```

Convergence

Convergence is the ability of a networking protocol to re-route around network outages. Static routing cannot do this. Dynamic routing protocols can all converge, but take various amounts of time to do this. Slow convergence can cause problems such as network loops which degrade network performance.

You may also hear robustness and redundancy used to describe networking protocols. In many ways they are the same thing as convergence. Robustness is the ability to keep working even though there are problems, including configuration problems as well as network outages. Redundancy involves having duplicate parts that can continue to function in the event of some malfunction, error, or outage. It is relatively easy to configure dynamic routing protocols to have backup routers and configurations that will continue to function no matter the network problem short of a total network failure.

IPv6 Support

IPv4 addressing is in common use everywhere around the world. IPv6 has much larger addresses and it is used by many large companies and government departments. IPv6 is not as common as IPv4 yet, but more companies are adopting it.

If your network uses IPv6, your dynamic routing protocol must support it. None of the dynamic routing protocols originally supported IPv6, but they all have additions, expansions, or new versions that do support IPv6. For more information, see “RIP and IPv6” on page 60, “BGP and IPv6” on page 99, “OSPFv3 and IPv6” on page 137, or “Integrated IS-IS” on page 184.

When to adopt dynamic routing

Static routing is more than enough to meet your networking needs when you have a small network. However, as your network grows, the question you need to answer is at what point do you adopt dynamic routing in your networking plan and start using it in your network? The main factors in this decision are typically:

- Budget
- Current network size and topology
- Expected network growth
- Available resources for ongoing maintenance

Budget

When making any business decision, the budget must always be considered. Static routing does not involve special hardware, fancy software, or expensive training courses.

Dynamic routing can include all of these extra expenses. Any new hardware, such as routers and switches, will need to support your chosen routing protocols. Network management software and routing protocol drivers may be necessary as well to help configure and maintain your more complex network. If the network administrators are not well versed in dynamic routing, either a training course or some hands-on learning time must be budgeted so they can administer the new network with confidence. Together, these factors can impact your budget.

Additionally, people will always account for network starting costs in the budgets, but usually leave out the ongoing cost of network maintenance. Any budget must provide for the hours that will be spent on updating the network routing equipment, and fixing any problems. Without that money in the budget, you may end up back at static routing before you know it.

Current network size and topology

As stated earlier static routing works well on small networks. At those networks get larger, routing takes longer, routing tables get very large, and general performance isn't what it could be.

Topology is a concern as well. If all your computers are in one building, its much easier to stay with static routing longer. However, connecting a number of locations will be easier with the move to dynamic routing.

If you have a network of 20 computers, you can still likely use static routing. If those computers are in two or three locations, static routing will still be a good choice for connecting them. Also, if you just connect to your ISP and don't worry about any special routing to do that, you are likely safe with just static routing.

If you have a network of 100 computers in one location, you can use static routing but it will be getting slower, more complex, and there won't be much room for expansion. If those 100 computers are spread across three or more locations, dynamic routing is the way to go.

If you have 1000 computers, you definitely need to use dynamic routing no matter how many locations you have.

Hopefully this section has given you an idea of what results you will likely experience from different sized networks using different routing protocols. Your choice of which dynamic routing protocol to use is partly determined by the network size, and topology.

Expected network growth

You may not be sure if your current network is ready for dynamic routing. However, if you are expecting rapid growth in the near future, it is a good idea to start planning for that growth now so you are ready for the coming expansion.

Static routing is very labor intensive. Each network device's routing table needs to be configured and maintained manually. If there is a large number of new computers being added to the network, they each need to have the static routing table configured and maintained. If devices are being moved around the network frequently, they must also be updated each time.

Instead, consider putting dynamic routing in place before those new computers are installed on the network. The installation issues can be worked out with a smaller and less complex network, and when those new computers or routers are added to the network there will be nowhere near the level of manual configuration required. Depending on the level of growth, this labor savings can be significant. For example, in an emergency you can drop a new router into a network or AS, wait for it to receive the routing updates from its neighbors, and then remove one of the neighbors. While the routes will not be the most effective possible, this method is much less work than static routing in the same situation, with less chance of mistakes.

Also, as your network grows and you add more routers, those new routers can help share the load in most dynamic routing configurations. For example if you have 4 OSPF routers and 20,000 external routes those few routers will be overwhelmed. But in a network with 15 OSPF routers they will better be able to handle that number of routes. Be aware though that adding more routers to your network will increase the amount of updates sent between the routers, which will use up a greater part of your bandwidth and use more bandwidth overall.

Available resources for ongoing maintenance

As touched on in the budget section, there must be resources dedicated to ongoing network maintenance, upgrades, and troubleshooting. These resources include administrator hours to configure and maintain the network, training for the administrator if needed, extra hardware and software as needed, and possible extra staff to help the administrator in emergencies. Without these resources, you will quickly find the network reverting to static routing out of necessity.

This is because:

- Routing software updates will require time.
- Routing hardware updates will require time.
- Office reorganizations or significant personnel movement will require time from a networking point of view.
- Networking problems that occur, such as failed hardware, require time to locate and fix the problem.

If the resources to accomplish these tasks are not budgeted, they will either not happen or not happen at the required level to continue operation. This will result in both the network administration staff and the network users being very frustrated.

A lack of maintenance budget will also result in increasingly heavy reliance on static routing as the network administrators are forced to use quick fixes for problems that come up. This invariably involves going to static routing, and dropping the more complex and time-consuming dynamic routing.

Choosing a routing protocol

One of that hardest decisions in routing can be choosing which routing protocol to use on your network. It can be easy to decide when static routing will not meet your needs, but how can you tell which dynamic routing protocol is best for your network and situation?

Here is a brief look at the routing protocols including their strongest and weakest points. The steps to choosing your routing protocol are:

1. Answer questions about your network
2. Dynamic routing terminology
3. Evaluate your chosen protocol
4. Implement your dynamic routing protocol

Answer questions about your network

Before you can decide what is best for your situation, you need to examine what the details of your situation are such as what you have for budget, equipment, and users.

The following questions will help you form a clear idea of your routing needs:

How many computers or devices are on your network?

It matters if you only have a few computers, or if you have many and if they are all at one location or not as well. All routing protocols can be run on any sized network, however it can be inefficient to run some on very small networks. However, routers and network hardware that support dynamic routing can be more expensive than more generic routers for static routing.

What applications typically run over the network?

Finding out what application your users are running will help you determine their needs and the needs of the network regarding bandwidth, quality of service, and other such issues.

What level of service do the users expect from the network?

Different network users have different expectations of the network. Its not critical for someone surfing the Internet to have 100% uptime, but it is required for a stock exchange network or a hospital.

Is there network expansion in your near future?

You may have a small network now, but if it will be growing quickly, you should plan for the expected size so you don't have to chance technologies again down the road.

What routing protocols do your networks connect to?

This is most often how routing protocol decisions are made. You need to be able to communicate easily with your service provider and neighbors, so often people simply use what everyone else is using.

Is security a major concern?

Some routing protocols have levels of authentication and other security features built in. Others do not. If security is important to you, be aware of this.

What is your budget — both initial and maintenance?

More robust and feature laden routing protocols generally mean more resources are required to keep them working well. Also more secure configurations require still more resources. This includes both set up costs, as well as ongoing maintenance costs. Ignore these costs at the risk of having to drop the adoption of the new routing protocol mid-change.

Evaluate your chosen protocol

Once you have examined the features of the routing protocols listed above and chosen the one that best meets your needs, you can set up an evaluation or test install of that protocol.

The test install is generally set up in a sandbox configuration so it will not affect critical network traffic. The aim of the test install is to prove that it will work on a larger scale on your network. So be sure that the test install mirrors your larger network well enough for you to discover any problems. If its too simplistic, these problems may not appear.

If your chosen protocol does not meet your goals choose a different protocol and repeat the evaluation process until either a protocol meets your needs, or you change your criteria.

Implement your dynamic routing protocol

You have examined your needs, selected the best matching dynamic routing protocol, tested it, and now you are ready to implement it with confidence.

This guide will help you configure your FortiGate unit to support your chosen dynamic routing protocol. Refer to the various sections in this guide as needed during your implementation to help ensure a smooth transition. Examples for each protocol have been included to show proper configurations for different types of networks.

Dynamic routing terminology

Dynamic routing is a complex subject. There are many routers on different networks and all can be configured differently. It become even more complicated when you add to this each routing protocol having slightly different names for similar features, and many configurable features for each protocol.

To better understand dynamic routing, here are some explanations of common dynamic routing terms.

- [Aggregated routes and addresses](#)
- [Autonomous system \(AS\)](#)
- [Area border router \(ABR\)](#)
- [Neighbor routers](#)
- [Route maps](#)
- [Access lists](#)
- [Bi-directional forwarding detection \(BFD\)](#)

For more details on a term as it applies to a dynamic routing protocol, see one of “[Border Gateway Protocol \(BGP\)](#)” on page 98, “[Routing Information Protocol \(RIP\)](#)” on page 59, or “[Open Shortest Path First \(OSPF\)](#)” on page 136.

Aggregated routes and addresses

Just as an aggregate interface combines multiple interfaces into one virtual interface, an aggregate route combines multiple routes into one. This reduces the amount of space those routes require in the routing tables of the routers along that route. The trade-off is a small amount of processing to aggregate and de-aggregate the routes at either end.

The benefit of this method is that you can combine many addresses into one, potentially reducing the routing table size immensely. The weakness of this method is if there are holes in the address range you are aggregating you need to decide if its better to break it into multiple ranges, or accept the possibility of failed routes to the missing addresses.

For information on aggregated routes in BGP, see “[ATOMIC_AGGREGATE](#)” on page 107, and “[Aggregate routes and addresses](#)” on page 110.

To manually aggregate the range of IP addresses from 192.168.1.100 to 192.168.1.103

1. Convert the addresses to binary

```
192.168.1.100 = 11000000 10101000 00000001 01100100
192.168.1.101 = 11000000 10101000 00000001 01100101
192.168.1.102 = 11000000 10101000 00000001 01100110
192.168.1.103 = 11000000 10101000 00000001 01100111
```

2. Determine the maximum number of matching bits common to the addresses.

There are 30-bits in common, with only the last 2-bits being different.

3 Record the common part of the address.

```
11000000 10101000 00000001 0110010X = 192.168.1.100
```

4 For the netmask, assume all the bits in the netmask are 1 except those that are different which are 0.

```
11111111 11111111 11111111 11111100 = 255.255.255.252
```

5 Combine the common address bits and the netmask.

```
192.168.1.100/255.255.255.252
```

Alternately the IP mask may be written as a single number:

```
192.168.1.100/2
```

6 As required, set variables and attributes to declare the routes have been aggregated, and what router did the aggregating.

Autonomous system (AS)

An Autonomous System (AS) is one or more connected networks that use the same routing protocol, and appear to be a single unit to any externally connected networks. For example an ISP may have a number of customer networks connected to it, but to any networks connected externally to the ISP it appears as one system or AS. An AS may also be referred to as a routing domain.

It should be noted that while OSPF routing takes place within one AS, the only part of OSPF that deals with the AS is the AS border router (ASBR).

There are multiple types of AS defined by how they are connected to other ASes. A multihomed AS is connected to at least two other ASes and has the benefit of redundancy — if one of those ASes goes down, your AS can still reach the Internet through its other connection. A stub AS only has one connection, and can be useful in specific configurations where limited access is desirable.

Each AS has a number assigned to it, known as an ASN. In an internal network, you can assign any ASN you like (a private AS number), but for networks connected to the Internet (public AS) you need to have an officially registered ASN from Internet Assigned Numbers Authority (IANA). ASNs from 1 to 64,511 are designated for public use.



As of January 2010, AS numbers are 4 bytes long instead of the former 2 bytes. RFC 4893 introduced 32-bit ASNs, which FortiGate units support for BGP and OSPF

Do you need your own AS?

The main factors in deciding if you need your own AS or if you should be part of someone else's are:

- exchanging external routing information
- many prefixes should exist in one AS as long as they use the same routing policy
- when you use a different routing protocol than your border gateway peers (for example your ISP uses BGP, and you use OSPF)
- connected to multiple other AS (multi-homed)

You should not create an AS for each prefix on your network. Neither should you be forced into an AS just so someone else can make AS-based policy decisions on your traffic.

There can be only one AS for any prefix on the Internet. This is to prevent routing issues.

What AS number to use?

In addition to overseeing IP address allocation and Domain Name Systems (DNS), the Internet Assigned Numbers Authority (IANA) assigns public AS numbers. The public AS numbers are from 1 to 64,511. The ASNs 0, 54272–64511, and 65535 are reserved by the IANA. These ASNs should not be used.

ASNs are assigned in blocks by the Internet Assigned Numbers Authority (IANA) to Regional Internet Registries (RIRs) who then assign ASNs to companies within that RIRs geographic area. Usually these companies are ISPs, and to receive an ASN you must complete the application process of the local RIR and be approved before being assigned an ASN. The RIRs names and regions are:

AFRINIC	Serves the African continent
APNIC	Asia-Pacific including China, India, and Japan
ARIN	American registry including Canada and United States
LACNIC	Latin America, including Mexico, Caribbean, Central and South America
RIPE NCC	Europe, the Middle East, former USSR, and parts of Central Asia

AS numbers from 64512 to 65534 are reserved for private use. Private AS numbers can be used for any internal networks with no outside connections to the Internet such as test networks, classroom labs, or other internal-only networks that do not access the outside world. You can also configure border routers to filter out any private ASNs before routing traffic to the outside world. If you must use private ASNs with public networks, this is the only way to configure them. However, it is risky because many other private networks could be using the same ASNs and conflicts will happen. It would be very much like your local 192.168.0.0 network being made public — the resulting problems would be widespread.

In 1996, when RFC 1930 was written only 5,100 ASes had been allocated and a little under 600 ASes were actively routed in the global Internet. Since that time many more public ASNs have been assigned, leaving only a small number. For this reason 32-bit ASNs (four-octet ASNs) were defined to provide more public ASNs. RFC 4893 defines 32-bit ASNs, and FortiGate units support these larger ASNs.

Area border router (ABR)

Routers within an AS advertise updates internally and only to each other. However, routers on the edge of the AS must communicate both with routers inside their AS and with routers external to their AS, often running a different routing protocol. These routers are called Area Border Routers (ABRs) or edge routers. Often ABRs run multiple routing protocols to be able to

redistribute traffic between different ASes that are running different protocols, such as the edge between an ISP's IS-IS routing network and a large company's OSPF network.

OSPF defines ABRs differently from other routers. In OSPF, an ABR is an OSPF router that connects another AS to the backbone AS, and is a member of all the areas it connects to. An OSPF ABR maintains a LSA database for each area that it is connected to. The concept of the edge router is present, but its the edge of the backbone instead of the edge of the OSPF supported ASes.

Neighbor routers

Routing involves routers communicating with each other. To do this, routers need to know information about each other. These routers are called neighbor routers, and are configured in each routing protocol. Each neighbor has custom settings since some routers may have functionality others routers lack. Neighbour routers are sometimes called peers.

Generally neighbor routers must be configured, and discovered by the rest of the network before they can be integrated to the routing calculations. This is a combination of the network administrator configuring the new router with its neighbor router addresses, and the routing network discovering the new router, such as the hello packets in OSPF. That discovery initiates communication between the new router and the rest of the network.

Route maps

Route maps are a way for the FortiGate unit to evaluate optimum routes for forwarding packets or suppressing the routing of packets to particular destinations. Compared to access lists, route maps support enhanced packet-matching criteria. In addition, route maps can be configured to permit or deny the addition of routes to the FortiGate unit routing table and make changes to routing information dynamically as defined through route-map rules.

Route maps can be used for limiting both received route updates, and sent route updates. This can include the redistribution of routes learned from other types of routing. For example if you don't want to advertise local static routes to external networks, you could use a route map to accomplish this.

The FortiGate unit compares the rules in a route map to the attributes of a route. The rules are examined in ascending order until one or more of the rules in the route map are found to match one or more of the route attributes.

As an administrator, route maps allow you to group a set of addresses together and assign them a meaningful name. Then during your configuration, you can use these route-maps to speed up configuration. The meaningful names ensure fewer mistakes during configuration as well.

The default rule in the route map (which the FortiGate unit applies last) denies all routes. For a route map to take effect, it must be called by a FortiGate unit routing process.

The syntax for route maps are:

```
config router route-map
  edit <route_map_name>
    set comments
    config rule
      edit <route_map_rule_id>
        set action
        set match-*
        set set-*
        ...
      end
    end
```


The `match-*` commands allow you to match various parts of a route. The `set-*` commands allow you to set routing information once a route is matched.

For an example of how route maps can be used to create receiving or sending “groups” in routing, see [“Redistributing and blocking routes in BGP” on page 129](#).

Access lists

Use this command to add, edit, or delete access lists. Access lists are filters used by FortiGate unit routing processes. For an access list to take effect, it must be called by a FortiGate unit routing process (for example, a process that supports RIP or OSPF). Use `access-list6` for IPv6 routing.

Access lists can be used to filter which updates are passed between routers, or which routes are redistributed to different networks and routing protocols. You can create lists of rules that will match all routes for a specific router or group of routers.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or to match the prefix and any more specific prefix.



If you are setting a prefix of 128.0.0.0, use the format 128.0.0.0/1. The default route, 0.0.0.0/0 can not be exactly matched with an access-list. A prefix-list must be used for this purpose.

The FortiGate unit attempts to match a packet against the rules in an access list starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found the default action is deny.

The syntax for access lists is:

```
config router access-list, access-list6
  edit <access_list_name>
    set comments
  config rule
    edit <access_list_id>
      set action
      set exact-match
      set prefix
      set prefix6
      set wildcard
```

For an example of how access lists can be used to create receiving or sending “groups” in routing, see [“Redistributing and blocking routes in BGP” on page 129](#).

Bi-directional forwarding detection (BFD)

Bi-directional Forwarding Detection (BFD) is a protocol used to quickly locate hardware failures in the network. Routers running BFD send packets to each other at a negotiated rate. If packets from a BFD-protected router fail to arrive, then that router is declared down. BFD communicates this information to the routing protocol and the routing information is updated.

BFD neighbors establish if BFD is enabled in OSPF or BFP routers that establish as neighbors.

The CLI commands associated with BFD include:

```
config router bgp
  config neighbor
    set bfd
```

```
config router ospf
  set bfd
```

Per-VDOM configuration:

```
config system settings
  set bfd
  set bfd-desired-min-tx
  set bfd-required-min-rx
  set bfd-detect-mult
  set bfd-dont-enforce-src-port
```

Per-interface (override) configuration:

```
config system interface
  edit <interface_name>
    set bfd enable
    set bfd-desired-min-tx
    set bfd-detect-mult
    set bfd-required-min-rx
```

For more information about BFD in BGP, see [“Bi-directional forwarding detection \(BFD\)”](#) on page 114.

IPv6 in dynamic routing

Unless otherwise stated, routing protocols apply to IPv4 addressing. This is the standard address format used. However, IPv6 is becoming more popular and new versions of the dynamic routing protocols have been introduced.

Dynamic routing supports IPv6 on your FortiGate unit. The new versions of these protocols and the corresponding RFCs are:

- **RIP next generation (RIPng)** — RFC 2080 - Routing Information Protocol next generation (RIPng). See [“RIP and IPv6”](#) on page 60.
- **BGP4+** — RFC 2545, and RFC 2858 Multiprotocol Extensions for IPv6 Inter-Domain Routing, and Multiprotocol Extensions for BGP-4 (MP-BGP) respectively. See [“BGP and IPv6”](#) on page 99.
- **OSPFv3** — RFC 2740 Open Shortest Path First version 3 (OSPFv3) for IPv6 support. See [“OSPFv3 and IPv6”](#) on page 137.
- **Integrated IS-IS** — RFC 5308 for IPv6 support. See [“Integrated IS-IS”](#) on page 184.

As with most advanced routing features on your FortiGate unit, IPv6 settings for dynamic routing protocols must be enabled before they will be visible in the GUI. To enable IPv6 configuration in the GUI, enable it in *System > Admin > Settings*. Alternatively, you can directly configure IPv6 for RIP, BGP, or OSPF protocols using CLI commands.

Routing Information Protocol (RIP)

This section describes the Routing Information Protocol (RIP).

The following topics are included in this section:

- [RIP background and concepts](#)
- [Troubleshooting RIP](#)
- [Simple RIP example](#)
- [RIPng — RIP and IPv6](#)

RIP background and concepts

This section contains:

- [Background](#)
- [Parts and terminology of RIP](#)
- [How RIP works](#)

Background

Routing Information Protocol (RIP) is a distance-vector routing protocol intended for small, relatively homogeneous networks. Its widespread use started when an early version of RIP was included with BSD v4.3 Linux as the routed daemon. The routing algorithm used by RIP, the Bellman–Ford algorithm, first saw widespread use as the initial routing algorithm of the ARPANET.

RIP benefits include being well suited to smaller networks, is in widespread use, near universal support on routing hardware, quick to configure, and works well if there are no redundant paths. However, RIP updates are sent out node-by-node so it can be slow to find a path around network outages. RIP also lacks good authentication, can not choose routes based on different quality of service methods, and can create network loops if you are not careful.

The FortiGate implementation of RIP supports RIP version 1 (see RFC 1058), RIP version 2 (see RFC 2453), and the IPv6 version RIPng (see RFC 2080).

RIP v1

In 1988 RIP version 1, defined in RFC 1058, was released. The RFC even states that RIP v1 is based on Linux routed due to it being a “defacto standard”.

It uses classful addressing and uses broadcasting to send out updates to router neighbors. There is no subnet information included in the routing updates in classful routing, and it does not support CIDR addressing — subnets must all be the same size. Also, route summarization is not possible.

RIP v1 has no router authentication method, so it is vulnerable to attacks through packet sniffing, and spoofing.

RIP v2

In 1993, RIP version 2 was developed to deal with the limitations of RIP v1. It was not standardized until 1998. This new version supports classless routing, and subnets of various sizes.

Router authentication was added in RIP v2 — it supports MD5. MD5 hashes are an older encryption method, but this is much improved over no security at all.

In RIP v2 the hop count limit remained at 15 to be backwards compatible with RIP v1.

RIP v2 uses multicasting to send the entire routing table to router neighbors, thereby reducing the traffic for devices that are not participating in RIP routing.

Routing tags were added as well, which allow internal routes or redistributed routes to be identified as such.

RIPng

RIPng, defined in RFC 2080, is an extension of RIP2 designed to support IPv6. However, RIPng varies from RIPv2 in that it is not fully backwards compatible with RIPv1.

- RIPng does not support RIPv1 update authentication, it relies on IPsec
- RIPng does not allow attaching tags to routes as in RIPv2
- RIPng requires specific encoding of the next hop for a set of route entries, unlike RIPv2 that encodes the next-hop into each route entry.

Parts and terminology of RIP

Before you can understand how RIP functions, you need to understand some of the main concepts and parts of RIP.

This section includes:

- [RIP and IPv6](#)
- [Default information originate option](#)
- [Garbage, timeout, and update timers](#)
- [Authentication and key-chain](#)
- [Access Lists](#)

RIP and IPv6

RIP Next Generation (RIPng) is a new version of RIP was released that includes support for IPv6.

The FortiGate unit command `config router ripng` is almost the same as `config router rip`, except that IPv6 addresses are used. Also if you are going to use prefix or access lists with RIPng, you must use the `config router access-list6` or `config prefix-list6` versions of those commands.

If you want to troubleshoot RIPng, it is the same as with RIP but specify the different protocol, and use IPv6 addresses. This applies to commands such as `get router info6` when you want to see the routing table, or other related information.

If you want to route IPv4 traffic over an IPv6 network, you can use the command `config system ip6-tunnel` to configure the FortiGate unit to do this. The IPv6 interface is configured under `config system interface`. All subnets between the source and destination addresses must support IPv6. This command is not supported in Transparent mode.

For example, you want to set up a tunnel on the port1 interface starting at 2002:C0A8:3201:: on your local network and tunnel it to address 2002:A0A:A01:: where it will need access to an IPv4 network again. Use the following command:

```
config system ipv6-tunnel
  edit test_tunnel
    set destination 2002:A0A:A01::
    set interface port1
    set source 2002:C0A8:3201::
  end
end
```

The CLI commands associated with RIPng include:

```
config router ripng
config router access-list6
config router prefix-list6
config system ipv6-tunnel
get router info6 *
```

Default information originate option

This is the second advanced option for RIP in the web-based manager, right after metric. Enabling default-information-originate will generate and advertise a default route into the FortiGate unit's RIP-enabled networks. The generated route may be based on routes learned through a dynamic routing protocol, routes in the routing table, or both. RIP does not create the default route unless you use the always option.

Select *Disable* if you experience any issues or if you wish to advertise your own static routes into RIP updates.

You can enable or disable default-information-originate in *Router > Dynamic > RIP*, under *Advanced Options*, or use the CLI.

The CLI commands associated with default information originate include:

```
config router rip
  set default-information-originate
end
```

Garbage, timeout, and update timers

RIP uses various timers to regulate its performance including a garbage timer, timeout timer, and update timer. The FortiGate unit default timer settings (30, 180, and 120 seconds respectively) are effective in most configurations — if you change these settings, ensure that the new settings are compatible with local routers and access servers.



The Timeout period should be at least three times longer than the Update period. If the Update timer is smaller than Timeout or Garbage timers, you will experience an error.

You can set the three RIP timers in *Router > Dynamic > RIP*, under *Advanced Options*, or use the CLI.

The CLI commands associated with garbage, timeout, and update timers include:

```
config router rip
  set garbage-timer
  set timeout-timer
  set update-timer
end
```

Garbage timer

The garbage timer is the amount of time (in seconds) that the FortiGate unit will advertise a route as being unreachable before deleting the route from the routing table. If this timer is shorter, it will keep more up-to-date routes in the routing table and remove old ones faster. This will result in a smaller routing table which is useful if you have a very large network, or if your network changes frequently.

Update timer

The update timer determines the interval between routing updates. Generally, this value is set to 30 seconds. There is some randomness added to help prevent network traffic congestion, which could result from all routers simultaneously attempting to update their neighbors. The update timer should be at least three times smaller than the timeout timer, otherwise you will experience an error.

If you are experiencing significant RIP traffic on your network, you can increase this interval to send fewer updates per minute. However, ensure you increase the interval for all the routers on your network or you will experience time outs that will degrade your network speed.

Timeout timer

The timeout timer is the maximum amount of time (in seconds) that a route is considered reachable while no updates are received for the route. This is the maximum time the FortiGate unit will keep a reachable route in the routing table while no updates for that route are received. If the FortiGate unit receives an update for the route before the timeout period expires, the timer is restarted. The timeout period should be at least three times longer than the depute period, otherwise you will experience an error.

If you are experiencing problems with routers not responding in time to updates, increase this timer. However, remember that longer timeout intervals result in longer overall update periods — it may be considerable time before the time the FortiGate unit is done waiting for all the timers to expire on unresponsive routes.

Authentication and key-chain

RIP version 2 uses authentication keys to ensure that the routing information exchanged between routers is reliable. RIP version 1 has no authentication. For authentication to work both the sending and receiving routers must be set to use authentication, and must be configured with the same keys.

The sending and receiving routers need to have their system dates and times synchronized to ensure both ends are using the same keys at the proper times. However, you can overlap the key lifetimes to ensure that a key is always available even if there is some difference in the system times.

A key chain is a list of one or more authentication keys including the send and receive lifetimes for each key. Keys are used for authenticating routing packets only during the specified lifetimes. The FortiGate unit migrates from one key to the next according to the scheduled send and receive lifetimes.

Key-chain is a CLI router command. You use this command to manage RIP version 2 authentication keys. You can add, edit or delete keys identified by the specified key number.

This example shows how to configure a key-chain with two keys that are valid sequentially in time. This example creates a key-chain called “rip_key” that has a password of “fortinet”. The accepted and send lifetimes are both set to the same values — a start time of 9:00am February 23, 2010 and an end time of 9:00am March 17, 2010. A second key is configured with a password of “my_fortigate” that is valid from March 17, 2010 9:01am to April 1 2010 9:00am. This “rip_key” keychain is then used on the port1 interface in RIP.

```
config router key-chain
  edit "rip_key"
    config key
      edit 1
        set accept-lifetime 09:00:00 23 02 2010 09:00:00 17 03 2010
        set key-string "fortinet"
        set send-lifetime 09:00:00 23 02 2010 09:00:00 17 03 2010
      next
      edit 2
        set accept-lifetime 09:01:00 17 03 2010 09:00:00 1 04 2010
        set key-string "my_fortigate"
        set send-lifetime 09:01:00 17 03 2010 09:00:00 1 04 2010
      next
    end
  end
config router rip
  config interface
    edit port1
      set auth-keychain "rip_key"
    end
  end
end
```

The CLI commands associated with authentication keys include:

```
config router key-chain

config router rip
  config interface
    edit <interface>
      set auth-keychain
      set auth-mode
      set auth-string
    end
  end
end
```

Access Lists

Access lists are filters used by FortiGate unit RIP and OSPF routing. An access list provides a list of IP addresses and the action to take for them — essentially an access list makes it easy to group addresses that will be treated the same into the same group, independent of their subnets or other matching qualities. You add a rule for each address or subnet that you want to include, specifying the action to take for it. For example if you wanted all traffic from one department to be routed a particular way, even in different buildings, you can add all the addresses to an access list and then handle that list all at once.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or to match the prefix and any more specific prefix.

The FortiGate unit attempts to match a packet against the rules in an access list starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found the default action is deny.

Access lists greatly speed up configuration and network management. When there is a problem, you can check each list instead of individual addresses. Also its easier to troubleshoot since if all addresses on one list have problems, it eliminates many possible causes right away.

If you are using the RIPng or OSPF+ IPv6 protocols you will need to use access-list6, the IPv6 version of access list. The only difference is that access-list6 uses IPv6 addresses.

For example, if you want to create an access list called `test_list` that only allows an exact match of `10.10.10.10` and `11.11.11.11`, enter the command:

```
config router access-list
  edit test_list
    config rule
      edit 1
        set prefix 10.10.10.10 255.255.255.255
        set action allow
        set exact-match enable
      next
      edit 2
        set prefix 11.11.11.11 255.255.255.255
        set action allow
        set exact-match enable
      end
    end
  end
```

Another example is if you want to deny ranges of addresses in IPv6 that start with the IPv6 equivalents of `10.10.10.10` and `11.11.11.11`, enter the command `access-list6` as follows:

```
config router access-list6
  edit test_list_ip6
    config rule
      edit 1
        set prefix6 2002:A0A:A0A:0:0:0:0:0/48
        set action deny
      next
      edit 2
        set prefix6 2002:B0B:B0B:0:0:0:0:0/48
        set action deny
      end
    end
  end
```

To use an `access_list`, you must call it from a routing protocol such as RIP. The following example uses the `access_list` from the earlier example called `test_list` to match routes coming in

on the port1 interface. When there is a match, it will add 3 to the hop count metric for those routes to artificially increase . Enter the following command:

```
config router rip
  config offset-list
    edit 5
      set access-list test_list
      set direction in
      set interface port1
      set offset 3
      set status enable
    end
```

If you are setting a prefix of 128.0.0.0, use the format 128.0.0.0/1. The default route, 0.0.0.0/0 can not be exactly matched with an access-list. A prefix-list must be used for this purpose

How RIP works

As one of the original modern dynamic routing protocols, RIP is straightforward. Its routing algorithm is not complex, there are some options to allow fine tuning, and it's relatively simple to configure RIP on FortiGate units.

From RFC 1058:

Distance vector algorithms are based on the exchange of only a small amount of information. Each entity (gateway or host) that participates in the routing protocol is assumed to keep information about all of the destinations within the system. Generally, information about all entities connected to one network is summarized by a single entry, which describes the route to all destinations on that network.

This section includes:

- [RIP versus static routing](#)
- [RIP metric – hop count](#)
- [The Bellman–Ford routing algorithm](#)
- [Passive versus active RIP interfaces](#)
- [RIP packet structure](#)

RIP versus static routing

RIP was one of the earliest dynamic routing protocols to work with IP addresses. As such, it is not as complex as more recent protocols. However, RIP is a big step forward from simple static routing.

While RIP may be slow in response to network outages, static routing has zero response. The same is true for convergence – static routing has zero convergence. Both RIP and static routing have the limited hop count, so its not a strength or a weakness. Count to infinity can be a problem, but typically can be fixed as it happens or is the result of a network outage that would cause even worse problems on static routing network.

This compares to static routing where each time a packet needs to be routed, the FortiGate unit can only send it to the next hop towards the destination. That next hop then forwards it, and so on until it arrives at its destination. RIP keeps more routing information on each router so your FortiGate unit can send the packet farther towards its destination before it has to be routed again towards its destination. RIP reduces the amount of table lookups and therefore fewer network resources than static routing. Also since RIP is updated on neighboring routes it is aware of new routes or dead routes that static routing would not be aware of.

Overall, RIP is a large step forward when compared to static routing.

RIP metric – hop count

RIP uses hop count as the metric for choosing the best route. A hop count of 1 represents a network that is connected directly to the FortiGate unit, while a hop count of 16 represents a network that cannot be reached. Each network that a packet travels through to reach its destination usually counts as one hop. When the FortiGate unit compares two routes to the same destination, it adds the route having the lowest hop count to the routing table. As you can see in [“RIP packet structure” on page 69](#), the hop count is part of a RIP v2 packet.

Similarly, when RIP is enabled on an interface, the FortiGate unit sends RIP responses to neighboring routers on a regular basis. The updates provide information about the routes in the FortiGate unit’s routing table, subject to the rules that you specify for advertising those routes. You can specify how often the FortiGate unit sends updates, the period of time a route can be kept in the routing table without being updated, and for routes that are not updated regularly you can specify the period of time that the unit advertises a route as unreachable before it is removed from the routing table.

If hops are weighted higher than one, it becomes very easy to reach the upper limit. This higher weighting will effectively limit the size of your network depending on the numbers used. Merely changing from the default of 1.0 to 1.5 will lower the effective hop count from 15 to 10. This is acceptable for smaller networks, but can be a problem as your network expands over time.

In RIP, you can use the offset command to artificially increase the hop count of a route. Doing this will make this route less preferred, and in turn it will get less traffic. Offsetting routes is useful when you have network connections of different bandwidths, different levels of reliability, or different costs. In each of these situations you still want the redundancy of multiple route access, but you don’t want the bulk of your traffic using these less preferred routes. For an example of RIP offset, see [“Access Lists” on page 63](#).

The Bellman–Ford routing algorithm

The routing algorithm used by RIP was first used in 1967 as the initial routing algorithm of the ARPANET. The Bellman–Ford algorithm is distributed because it involves a number of nodes (routers) within an Autonomous system, and consists of the following steps:

1. Each node calculates the distances between itself and all other nodes within the AS and stores this information as a table.
2. Each node sends its table to all neighboring nodes.
3. When a node receives distance tables from its neighbors, it calculates the shortest routes to all other nodes and updates its own table to reflect any changes.

To examine how this algorithm functions let’s look at a network with 4 routers — routers 1 through 4. The distance from router1 to router2 is 2 hops, 1 to 3 is 3 hops, and 2 to 3 is 4 hops. Router4 is only connected to routers 2 and 3, each distance being 2 hops.

1. Router1 finds all the distance to the other three routers — router 2 is 2, router 3 is 3. Router1 doesn’t have a route to router 4.
2. Routers 2 through 4 do the same calculations from their point of views.
3. Once router 1 gets an update from router 2 or 3, it will get their route to router 4. At that point it now has a route to router 4 and installs that in its local table.
4. If router1 gets an update from router3 first, it has a hop count of 5 to reach router4. But when router2 sends its update, router1 will go with router2’s shorter 4 hops to reach router4. Future updates don’t change this unless they are shorter than 4 hops, or the routing table route goes down.

Figure 3: RIP algorithm example in 4 steps

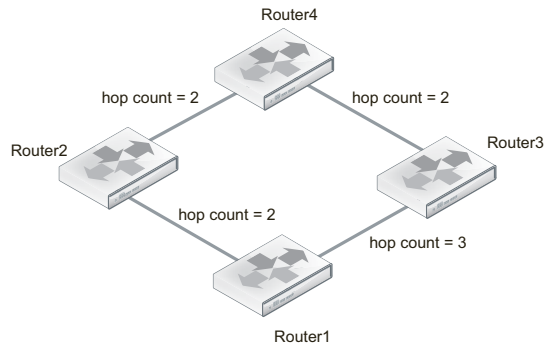
Step 1

Router1 finds the distance to other routers in the network.

It currently has no route to Router4.

Router1 routing table:

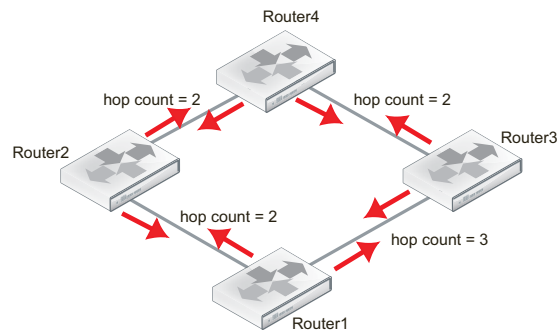
- Distance to Router2 = 2 hops.
- Distance to Router3 = 3 hops.



Step 2

All routers do the same as Router1, and send out updates containing their routing table.

Note that Router1 and Router4 do not update each other, but rely on Router2 and Router3 to pass along accurate updates.




Step 3

Each router looks at the updates it has received, and adds any new or shorter routes to its table.

Router1 updated table:

- Distance to Router2 = 2 hops.
- Distance to Router3 = 3 hops.
- Distance to Router4 = 4 or 5 hops.

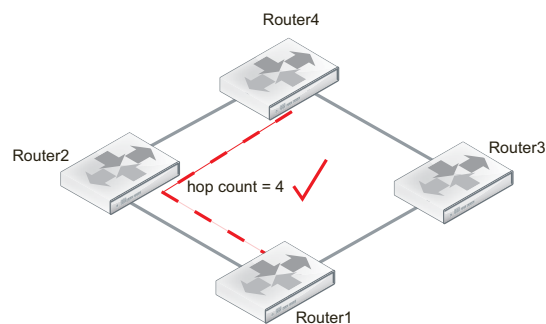


Step 4

Router1 installs the shortest route to Router4, and the other routes to it are removed from the routing table.

Router1 complete table:

- Distance to Router2 = 2 hops.
- Distance to Router3 = 3 hops.
- Distance to Router4 = 4 hops.



The good part about the Bellman-Ford algorithm in RIP is that the router only uses the information it needs from the update. If there are no newer, better routes than the ones the router already has in its routing table, there is no need to change its routing table. And no change means no additional update, so less traffic. But even when there is update traffic, the RIP packets are very small so it takes many updates to affect overall network bandwidth. For more information about RIP packets, see [“RIP packet structure” on page 69](#).

The main disadvantage of the Bellman–Ford algorithm in RIP is that it doesn't take weightings into consideration. While it is possible to assign different weights to routes in RIP, doing so severely limits the effective network size by reducing the hop count limit. Also other dynamic routing protocols can take route qualities, such as reliability or delay, into consideration to provide not only the physically shortest but also the fastest or more reliable routes as you choose.

Another disadvantage of the Bellman-Ford algorithm is due to the slow updates passed from one RIP router to the next. This results in a slow response to changes in the network topology, which in turn results in more attempts to use routes that are down, which wastes time and network resources.

Passive versus active RIP interfaces

Normally the FortiGate unit's routing table is kept up to date by periodically asking the neighbors for routes, and sending your routing updates out. This has the downside of generating a lot of extra traffic for large networks. The solution to this problem is passive interfaces.

An standard interface that supports RIP is active by default — it both sends and receives updates by actively communicating with its neighbors. A passive RIP interface does not send out updates — it just listens to the updates of other routers. This is useful in reducing network traffic, and if there are redundant routers in the network that would be sending out essentially the same updates all the time.

The following example shows how to create a passive RIP v2 interface on port1, using MD5 authentication and a key-chain called `passiveRIPv2` that has already been configured. Note that in the CLI, you enable passive by disabling `send-version2-broadcast`.

To create a passive RIP interface - web-based manager

1. Go to *Router > Dynamic > RIP*.
2. Under *Interfaces*, select *Create*.
3. Select port1 as the *Interface*.
4. Select 2 as both the *Send Version* and *Receive Version*.
5. Select MD5 for *Authentication*.
6. Select the `passiveRIPv2` *Key-chain*.
7. Select *Passive Interface*.
8. Select *OK* to accept this configuration, and return to the main RIP display page.

To create a passive RIP v2 interface on port1 using MD5 authentication- CLI

```
config router rip
  config interface
    edit port1
      set send-version2-broadcast disable
      set auth-keychain "passiveRIPv2"
      set auth-mode md5
      set receive-version 2
      set send-version 2
    end
  end
```

RIP packet structure

It is hard to fully understand a routing protocol without knowing what information is carried in its packets. Knowing what information is exchanged between routers and how will help you better understand the RIP protocol, and better configure your network for it.

This section provides information on the contents of RIP 1 and RIP 2 packets.

RIP version 1

RIP version 1, or RIP IP packets are 24 bytes in length, with some empty areas left for future expansion.

Table 7: RIP IP packets

1-byte command	1-byte version	2-byte zero field	2-byte AFI	2-byte zero field
4-byte IP address	4-byte zero field	4-byte zero field	4-byte metric	

A RIP 1 packet contains the following fields:

Command — Indicates whether the packet is a request or a response. The request asks that a router send all or part of its routing table. The response can be an unsolicited regular routing update or a reply to a request. Responses contain routing table entries. Multiple RIP packets are used to convey information from large routing tables.

Version — Specifies the RIP version used. This field can signal different potentially incompatible versions.

Zero field — This field defaults to zero, and is not used by RFC 1058 RIP.

Address-family identifier (AFI) — Specifies the address family used. RIP is designed to carry routing information for several different protocols. Each entry has an address-family identifier to indicate the type of address being specified. The AFI for IP is 2.

IP Address — Specifies the IP address for the entry.

Metric — This is the number of hops or routers traversed along the route on its trip to the destination. The metric is between 1 and 15 for that number of hops. If the route is unreachable the metric is 16.

RIP version 2

RIP version 2 has more features than RIP 1, which is reflected in its packets which carry more information. All but one of the empty zero fields in RIP 1 packets are used in RIP 2.

Table 8: RIP 2 packets

1-byte command	1-byte version	2-byte unused	2-byte AFI	2-byte route tag	4-byte IP address	4-byte subnet	4-byte next hop	4-byte metric
----------------	----------------	---------------	------------	------------------	-------------------	---------------	-----------------	---------------

A RIP 2 packet contains fields described above in RIP 1, as well as the following:

Unused — Has a value set to zero, and is intended for future use

Route tag — Provides a method for distinguishing between internal routes learned by RIP and external routes learned from other protocols.

Subnet mask — Contains the subnet mask for the entry. If this field is zero, no subnet mask has been specified for the entry.

Next hop — Indicates the IP address of the next hop to which packets for the entry should be forwarded.

Troubleshooting RIP

This section is about troubleshooting RIP. For general troubleshooting information, see the FortiOS Handbook Troubleshooting chapter.

This section includes:

- [Routing Loops](#)
- [Holddowns and Triggers for updates](#)
- [Split horizon and Poison reverse updates](#)
- [Debugging IPv6 on RIPng](#)

Routing Loops

Normally in routing, a path between two addresses is chosen and traffic is routed along that path from one address to the other. When there is a routing loop, that normal path doubles back on itself creating a loop. When there are loops, the network has problems getting information to its destination and also prevents it from returning to the source to report the inaccessible destination.

A routing loop happens when a normally functioning network has an outage, and one or more routers are offline. When packets encounter this, an alternate route is attempted to maneuver around the outage. During this phase it is possible for a route to be attempted that involves going back a hop, and trying a different hop forward. If that hop forward is blocked by the outage as well, a hop back and possibly the original hop forward may be selected. You can see if this continues, how it can consume not only network bandwidth but also many resources on those routers affected. The worst part is this situation will continue until the network administrator changes the router settings, or the downed routers come back online.

Routing loops' effect on the network

In addition to this “traffic jam” of routed packets, every time the routing table for a router changes that router sends an update out to all of the RIP routers connected to it. In a network loop, its possible for a router to change its routes very quickly as it tries and fails along these new routes. This can quickly result in a flood of updates being sent out, which can effectively grind the network to a halt until the problem is fixed.

How can you spot a routing loop

Any time network traffic slows down, you will be asking yourself if it is a network loop or not. Often slowdowns are normal, they are not a full stoppage, and normal traffic resumes in a short period of time.

If the slow down is a full halt of traffic or a major slowdown does not return to normal quickly, you need to do serious troubleshooting quickly.

If you aren't running SNMP, dead gateway detection, or you have non-Fortinet routers in your network, you can use networking tools such as ping and traceroute to define the outage on your network and begin to fix it. Ping, traceroute, and other basic troubleshooting tools are largely the same between static and dynamic, and are covered in [“Troubleshooting static routing” on page 21](#).

Check your logs

If your routers log events to a central location, it can be easy to check the logs for your network for any outages.

On your FortiGate unit, go to *Log & Report*. You will want to look at both event logs and traffic logs. Events to look for will generally fall under CPU and memory usage, interfaces going offline (due to dead gateway detection), and other similar system events.

Once you have found and fixed your network problem, you can go back to the logs and create a report to better see how things developed during the problem. This type of forensics analysis can better help you prepare for next time.

Use SNMP network monitoring

If your network had no problems one minute and slows to a halt the next, chances are something changed to cause that problem. Most of the time an offline router is the cause, and once you find that router and bring it back online, things will return to normal.

If you can enable a hardware monitoring system such as SNMP or sFlow on your routers, you can be notified of the outage and where it is exactly as soon as it happens.

Ideally you can configure SNMP on all your FortiGate routers and be alerted to all outages as they occur.

To use SNMP to detect potential routing loops

1. Go to *System > Config > SNMP*.

2. Enable *SMTP Agent* and select *Apply*.

Optionally enter the *Description*, *Location*, and *Contact* information for this device for easier location of the problem report.

3. Under *SNMP v1/v2* or *SNMP v3* as appropriate, select *Create New*.

SNMP v3

User Name	Enter the SNMP user ID.
Security Level	Select authentication or privacy as desired. Select the authentication or privacy algorithms to use and enter the required passwords.
Notification Host	Enter the IP addresses of up to 16 hosts to notify.
Enable Query	Select. The <i>Port</i> should be 161. Ensure that your security policies allow ports 161 and 162 (SNMP queries and traps) to pass.

SNMP v1/v2

Hosts	Enter the IP addresses of up to 8 hosts to notify. You can also specify the network <i>Interface</i> , or leave it as <i>ANY</i> .
--------------	--

Queries	Enable v1 and/or v2 as needed. The <i>Port</i> should be 161. Ensure that your security policies allow port 161 to pass.
----------------	--

Traps	Enable v1 and/or v2 as needed. The <i>Port</i> should be 162. Ensure that your security policies allow port 162 to pass.
--------------	--

4. Select the events for which you want notification. For routing loops this should include *CPU usage is high*, *Memory is low*, and possibly *Log disk space is low*. If there are problems the log will be filling up quickly, and the FortiGate unit's resources will be overused.
5. Configure SNMP host (manager) software on your administration computer. This will monitor the SNMP information sent out by the FortiGate unit. Typically you can configure this software to alert you to outages or CPU spikes that may indicate a routing loop.

Use dead gateway detection and e-mail alerts

Another tool available to you on FortiGate units is the dead gateway detection. This feature allows the FortiGate unit to ping a gateway at regular intervals to ensure it is online and working. When the gateway is not accessible, that interface is marked as down.

To detect possible routing loops with dead gateway detection and e-mail alerts

1. To configure dead gateway detection, go to *Router > Static > Settings* and select *Create New*.
2. Enter the *Ping Server* IP address and select the *Interface* that connects to it.
3. Set the *Ping Interval* (how often to send a ping), and *Failover Threshold* (how many lost pings is considered a failure). A smaller interval and smaller number of lost pings will result in faster detection, but will create more traffic on your network.

To configure notification of failed gateways

1. Go to *Log&Report > Log Config > Alert E-mail*.
2. Enter your email details.
3. Select the *Configuration changes* event.
4. Select *Apply*.

You might also want to log CPU and Memory usage as a network outage will cause your CPU activity to spike.



If you have VDOMs configured, you will have to enter the basic SMTP server information in the Global section, and the rest of the configuration within the VDOM that includes this interface.

After this configuration, when this interface on the FortiGate unit cannot connect to the next router, the FortiGate unit will bring down the interface and alert you with an email about the outage.

Look at the packet flow

If you want to see what is happening on your network, look at the packets travelling on the network. This is same idea as police pulling over a car and asking the driver where they have been, and what the conditions were like.

The method used in the troubleshooting sections “[Debugging IPv6 on RIPng](#)” on page 74 and on debugging the packet flow apply here as well. In this situation, you are looking for routes that have metrics higher than 15 as that indicates they are unreachable.

Ideally if you debug the flow of the packets, and record the routes that are unreachable, you can create an accurate picture of the network outage.

Action to take on discovering a routing loop

Once you have mapped the problem on your network, and determined it is in fact a routing loop there are a number of steps to take in correcting it.

1. Get any offline routers back online. This may be a simple reboot, or you may have to replace hardware. Often this first step will restore your network to its normal operation, once the routing tables finish being updated.
2. Change your routing configuration on the edges of the outage. Even if step 1 brought your network back online, you should consider making changes to improve your network before the next outage occurs. These changes can include configuring features like holddowns and triggers for updates, split horizon, and poison reverse updates.

Holddowns and Triggers for updates

One of the potential problems with RIP is the frequent routing table updates that are sent every time there is a change to the routing table. If your network has many RIP routers, these updates can start to slow down your network. Also if you have a particular route that has bad hardware, it might be going up and down frequently, which will generate an overload of routing table updates.

One of the most common solutions to this problem is to use holddown timers and triggers for updates. These slow down the updates that are sent out, and help prevent a potential flood.

Holddown Timers

The holddown timer activates when a route is marked down. Until the timer expires, the router does not accept any new information about that route. This is very useful if you have a flapping route because it will prevent your router from sending out updates and being part of the problem in flooding the network. The potential down side is if the route comes back up while the timer has not expired, that route will be unavailable for that period of time. This is only a problem if this is a major route used by the majority of your traffic. Otherwise, this is a minor problem as traffic can be re-routed around the outage.

Triggers

Triggered RIP is an alternate update structure that is based around limiting updates to only specific circumstances. The most basic difference is that the routing table will only be updated when a specific request is sent to update, as opposed to every time the routing table changes. Updates are also triggered when a unit is 'powered on', which can include addition of new interfaces or devices to the routing structure, or devices returning to being available after being unreachable.

Split horizon and Poison reverse updates

Split horizon is best explained with an example. You have three routers linked serially, let's call them A, B, and C. A is only linked to B, C is only linked to B, and B is linked to both A and C. To get to C, A must go through B. If the link to C goes down, it is possible that B will try to use A's route to get to C. This route is A-B-C, so it will loop endlessly between A and B.

This situation is called a split horizon because from B's point of view the horizon stretches out in each direction, but in reality it only is on one side.

Poison reverse is the method used to prevent routes from running into split horizon problems. Poison reverse "poisons" routes away from the destination that use the current router in their

route to the destination. This “poisoned” route is marked as unreachable for routers that cannot use it. In RIP this means that route is marked with a distance of 16.

Debugging IPv6 on RIPng

The debug commands are very useful to see what is happening on the network at the packet level. There are a few changes to debugging the packet flow when debugging IPv6.

The following CLI commands specify both IPv6 and RIP, so only RIPng packets will be reported. The output from these commands will show you the RIPng traffic on your FortiGate unit including RECV, SEND, and UPDATE actions.

The addresses are in IPv6 format.

```
diagnose debug enable
diagnose ipv6 router rip level info
diagnose ipv6 router rip all enable
```

These three commands will:

- turn on debugging in general
- set the debug level to information, a verbose reporting level
- turn on all rip router settings

Part of the information displayed from the debugging is the metric (hop count). If the metric is 16, then that destination is unreachable since the maximum hop count is 15.

In general, you should see an update announcement, followed by the routing table being sent out, and a received reply in response.

For more information, see [“Testing the IPv6 RIPng information”](#) on page 97

Simple RIP example

This is an example of a typical medium sized network configuration using RIP routing.

Your company has 3 small local networks, one for each department. These networks are connected by RIP, and then connected to the Internet. Each subnet has more than one route, for redundancy. There are two central routers that are both connected to the internet, and to the other networks. If one of those routers goes down, the whole network can continue to function normally.

The ISP is running RIP, so no importing or exporting routes is required on the side of the network. However, since the internal networks have static networking running those will need to be redistributed through the RIP network.

To keep the example simple, there will be no authentication of router traffic.

With RIP properly configured, if the device fails or temporarily goes offline, the routes will change and traffic will continue to flow. RIP is good for a smaller network due to its lack of complex configurations.

This section includes the following topics:

- [Network layout and assumptions](#)
- [General configuration steps](#)
- [Configuring the FortiGate units system information](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

Network layout and assumptions

Basic network layout

Your company has 3 departments each with their own network — Sales, R&D, and Accounting. Each network has routers that are not running RIP as well as FortiGate units running RIP.

The R&D network has two RIP routers, and each is connected to both other departments as well as being connected to the Internet through the ISP router. The links to the Internet are indicated in black.

The three internal networks do not run RIP. They use static routing because they are small networks. This means the FortiGate units have to redistribute any static routes they learn so that the internal networks can communicate with each other.


Where possible in this example, the default values will be used or the most general settings. This is intended to provide an easier configuration that will require less troubleshooting.

In this example the routers, networks, interfaces used, and IP addresses are as follows. Note that the Interfaces that connect Router2 and Router3 also connect to the R&D network.

Table 9: Rip example network topology

Network	Router	Interface & Alias	IP address
Sales	Router1	port1 (internal)	10.11.101.101
		port2 (router2)	10.11.201.101
		port3 (router3)	10.11.202.101
R&D	Router2	port1 (internal)	10.12.101.102
		port2 (router1)	10.11.201.102
		port3 (router4)	10.14.201.102
		port4 (ISP)	172.20.120.102
	Router3	port1 (internal)	10.12.101.103
		port2 (router1)	10.11.201.103
		port3 (router4)	10.14.202.103
		port4 (ISP)	172.20.120.103
Accounting	Router4	port1 (internal)	10.14.101.104
		port2 (router2)	10.14.201.104
		port3 (router3)	10.14.202.104

Figure 4: Network topology for the simple RIP example



Assumptions

The following assumptions have been made concerning this example.

- All FortiGate units have 5.0 firmware, and are running factory default settings.
- All CLI and web-based manager navigation assumes the unit is running in NAT/Route operating mode, with VDOMs disabled.
- All FortiGate units have interfaces labelled port1 through port4 as required.
- All firewalls have been configured for each FortiGate unit to allow the required traffic to flow across interfaces.
- Only FortiGate units are running RIP on the internal networks.
- Router2 and Router3 are connected through the internal network for R&D.
- Router2 and Router3 each have their own connection to the Internet, indicated in black in Figure 4.

General configuration steps

This example is very straight forward. The only steps involved are:

- [Configuring the FortiGate units system information](#)
- [Configuring FortiGate unit RIP router information](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

Configuring the FortiGate units system information

Each FortiGate unit needs their hostname, and interfaces configured.

For IP numbering, Router2 and Router3 use the other routers numbering where needed.

Router2 and Router3 have dead gateway detection enabled on the ISP interfaces using Ping. Remember to contact the ISP and confirm their server has ping enabled.

Configure the hostname, interfaces, and default route

To configure Router1 system information - web-based manager

1. Go to *System > Dashboard > Status > System Information*.
2. Next to *Host Name* select *Change*, and enter "Router1".
3. Go to *Router > Static > Static Routes*.
4. Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port2 (router2)
Gateway	172.20.120.5/255.255.255.0
Distance	40

5. Enter a second default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port3 (router3)
Gateway	172.20.120.5/255.255.255.0
Distance	40

6. Go to *System > Network > Interfaces*.
7. Edit port1 (internal) interface.
8. Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.11.101.101/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal sales network
Administrative Status	Up

9. Edit port2 (router2) interface.
10. Set the following information, and select *OK*.

Alias	router2
IP/Netmask	10.11.201.101/255.255.255.0
Administrative Access	HTTPS SSH PING

Description	Link to R&D network & internet through Router2
Administrative Status	Up

11. Edit port3 (router3) interface.

12. Set the following information, and select *OK*.

Alias	router3
IP/Netmask	10.11.202.101/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to R&D network and internet through Router3
Administrative Status	Up

To configure Router1 system information - CLI

```
config system global
  set hostname Router1
end
```

```

config router static
  edit 1
    set device "port2"
    set distance 45
    set gateway 10.11.201.102
  next
  edit 2
    set device "port3"
    set distance 45
    set gateway 10.11.202.103
  end
end

config system interface
  edit port1
    set alias internal
    set ip 10.11.101.101/255.255.255.0
    set allowaccess https ssh ping
    set description "Internal sales network"
  next
  edit port2
    set alias ISP
    set allowaccess https ssh ping
    set ip 10.11.201.101/255.255.255.0
    set description "Link to R&D network & internet through Router2"
  next
  edit port3
    set alias router3
    set ip 10.11.202.101/255.255.255.0
    set allowaccess https ssh ping
    set description "Link to R&D network & internet through Router2"
  end
end

```

To configure Router2 system information - web-based manager

1. Go to *System > Dashboard > Status > System Information*.
2. Next to *Host Name* select *Change*, and enter "Router2".
3. Go to *Router > Static > Static Routes*.
4. Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port4 (ISP)
Gateway	172.20.120.5/255.255.255.0
Distance	5

5. Go to *System > Network > Interfaces*.
6. Edit port1 (internal) interface.
7. Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.12.101.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	R&D internal network and Router3
Administrative Status	Up

8. Edit port2 (router1) interface.
9. Set the following information, and select *OK*.

Alias	router1
IP/Netmask	10.12.201.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to Router1 and the Sales network
Administrative Status	Up

10. Edit port3 (router4) interface.
11. Set the following information, and select *OK*.

Alias	router4
IP/Netmask	10.12.301.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to Router4 and the accounting network
Administrative Status	Up

12. Edit port4 (ISP) interface.
13. Set the following information, and select *OK*.

Alias	ISP
IP/Netmask	172.20.120.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Detect Interface Status for Gateway Load Balancing	enable
Detect Server	172.20.120.5
Detect Protocol	Ping

Detect Interface Status for Gateway Load Balancing	enable
Description	Internet through ISP
Administrative Status	Up

To configure Router2 system information - CLI

```

config system global
    set hostname Router2
end
config router static
    edit 1
        set device "port4"
        set distance 5
        set gateway 172.20.130.5
    end
end
config system interface
    edit port1
        set alias internal
        set ip 10.11.101.102/255.255.255.0
        set allowaccess https ssh ping
        set description "Internal RnD network and Router3"
    next
    edit port2
        set alias router1
        set allowaccess https ssh ping
        set ip 10.11.201.102/255.255.255.0
        set description "Link to Router1"
    next
    edit port3
        set alias router3
        set ip 10.14.202.102/255.255.255.0
        set allowaccess https ssh ping
        set description "Link to Router4"
    next
    edit port4
        set alias ISP
        set ip 172.20.120.102/255.255.255.0
        set allowaccess https ssh ping
        set description "ISP and internet"
    end
end
end

```

To configure Router3 system information - web-based manager

1. Go to *System > Dashboard > Status > System Information*.
2. Next to *Host Name* select *Change*, and enter "Router3".
3. Go to *Router > Static > Static Route*.

4. Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port4 (ISP)
Gateway	172.20.120.5/255.255.255.0
Distance	5

5. Go to *System > Network > Interfaces*.
6. Edit port1 (internal) interface.
7. Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.12.101.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	R&D internal network and Router2
Administrative Status	Up

8. Edit port2 (router1) interface.
9. Set the following information, and select *OK*.

Alias	router1
IP/Netmask	10.13.201.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to Router1 and Sales network
Administrative Status	Up

10. Edit port3 (router4) interface.
11. Set the following information, and select *OK*.

Alias	router4
IP/Netmask	10.13.301.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to Router4 and accounting network
Administrative Status	Up

12. Edit port4 (ISP) interface.

13. Set the following information, and select *OK*.

Alias	ISP
IP/Netmask	172.20.120.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Detect Interface Status for Gateway Load Balancing	enable
Detect Server	172.20.120.5
Detect Protocol	Ping
Description	Internet and ISP
Administrative Status	Up

To configure Router3 system information - CLI

```
config system global
  set hostname Router3
end

config router static
  edit 1
    set device "port4"
    set distance 5
    set gateway 172.20.130.5
  end
end

config system interface
  edit port1
    set alias internal
    set ip 10.12.101.103/255.255.255.0
    set allowaccess https ssh ping
    set description "Internal RnD network and Router2"
  next
  edit port2
    set alias ISP
    set allowaccess https ssh ping
    set ip 10.11.201.103/255.255.255.0
    set description "Link to Router1"
  next
  edit port3
    set alias router3
    set ip 10.14.202.103/255.255.255.0
    set allowaccess https ssh ping
    set description "Link to Router4"
  next
  edit port4
    set alias ISP
    set ip 172.20.120.103/255.255.255.0
    set allowaccess https ssh ping
    set description "ISP and internet"
  end
end
```

To configure Router4 system information - web-based manager

1. Go to *System > Dashboard > Status > System Information*.
2. Next to *Host Name* select *Change*, and enter "Router4".
3. Go to *Router > Static > Static Routes*.
4. Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port2 (router2)

Gateway	172.20.120.5/255.255.255.0
Distance	40

5. Enter a second default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port3 (router3)
Gateway	172.20.120.5/255.255.255.0
Distance	40

6. Go to *System > Network > Interfaces*.
 7. Edit port 1 (internal) interface.
 8. Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.14.101.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal accounting network
Administrative Status	Up

9. Edit port 2 (router2) interface.
 10. Set the following information, and select *OK*.

Alias	router2
IP/Netmask	10.14.201.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to R&D network & internet through Router2
Administrative Status	Up

11. Edit port 3 (router3) interface.
 12. Set the following information, and select *OK*.

Alias	router3
IP/Netmask	10.14.301.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Link to R&D network and internet through Router3
Administrative Status	Up

To configure Router4 system information - CLI

```
config system global
  set hostname Router4
end

config router static
  edit 1
    set device "port2"
    set distance 45
    set gateway 10.14.201.102
  next
  edit 2
    set device "port3"
    set distance 45
    set gateway 10.14.202.103
  end
end

config system interface
  edit port1
    set alias internal
    set ip 10.14.101.104/255.255.255.0
    set allowaccess https ssh ping
    set description "Internal sales network"
  next
  edit port2
    set alias router2
    set allowaccess https ssh ping
    set ip 10.14.201.104/255.255.255.0
    set description "Link to R&D network & internet through Router2"
  next
  edit port3
    set alias router3
    set ip 10.14.202.104/255.255.255.0
    set allowaccess https ssh ping
    set description "Link to R&D network & internet through Router2"
  end
end
```

Configuring FortiGate unit RIP router information

With the interfaces configured, RIP can now be configured on the FortiGate units.

For each FortiGate unit the following steps will be taken:

- configure RIP version used
- redistribute static networks
- add networks serviced by RIP
- add interfaces that support RIP on the FortiGate unit

Router1 and Router4 are configured the same. Router2 and Router3 are configured the same. These routers will be grouped accordingly for the following procedures — repeat the procedures once for each FortiGate unit.

Configure RIP settings on Router1 and Router4 - web-based manager

1. Go to *Router > Dynamic > RIP*.
2. Select 2 for *RIP Version*.
3. In *Advanced Options*, under *Redistribute* enable *Static*.
4. Leave the other *Advanced Options* at default values.
5. Enter the following networks, and select *Add* after each:
 - 10.11.0.0/255.255.0.0
 - 10.12.0.0/255.255.0.0
 - 10.14.0.0/255.255.0.0
 - 172.20.120.0/255.255.255.0
6. For interface, select *Create New* and set the following information.

Interface	port1 (internal)
Send Version	Both
Receive Version	Both
Authentication	None
Passive Interface	disabled

7. For interface, select *Create New* and set the following information.

Interface	port2 (router2)
Send Version	Both
Receive Version	Both
Authentication	None
Passive Interface	disabled

8. For interface, select *Create New* and set the following information.

Interface	port3 (router3)
Send Version	Both
Receive Version	Both
Authentication	None
Passive Interface	disabled

Configure RIP settings on Router1 and Router4 - CLI

```
config router rip
  set version 2
  config interface
    edit "port1"
      set receive-version 1 2
      set send-version 1 2
    next
    edit "port2"
      set receive-version 1 2
      set send-version 1 2
    next
    edit "port3"
      set receive-version 1 2
      set send-version 1 2
    end
  config network
    edit 1
      set prefix 10.11.0.0 255.255.0.0
    next
    edit 2
      set prefix 10.12.0.0 255.255.0.0
    next
    edit 3
      set prefix 10.14.0.0 255.255.0.0
    next
    edit 4
      set prefix 172.20.120.0 255.255.255.0
    end
  config redistribute "static"
    set status enable
  end
end
```

Configure RIP settings on Router2 and Router3- web-based manager

1. Go to *Router > Dynamic > RIP*.
2. Select 2 for *RIP Version*.
3. In *Advanced Options*, under *Redistribute* enable *Static*.
4. Leave the other *Advanced Options* at default values.
5. Enter the following networks, and select *Add* after each:
 - 10.11.0.0/255.255.0.0
 - 10.12.0.0/255.255.0.0
 - 10.14.0.0/255.255.0.0
 - 172.20.120.0/255.255.255.0

6. For interface, select *Create New* and set the following information.

Interface	port1 (internal)
Send Version	Both
Receive Version	Both
Authentication	None
Passive Interface	disabled

7. For interface, select *Create New* and set the following information.

Interface	port2 (router1)
Send Version	Both
Receive Version	Both
Authentication	None
Passive Interface	disabled

8. For interface, select *Create New* and set the following information.

Interface	port3 (router4)
Send Version	Both
Receive Version	Both
Authentication	None
Passive Interface	disabled

9. For interface, select *Create New* and set the following information.

Interface	port4 (ISP)
Send Version	Both
Receive Version	Both
Authentication	None
Passive Interface	disabled

Configure RIP settings on Router2 and Router3- web-based manager

```
config router rip
  set version 2
  config interface
    edit "port1"
      set receive-version 1 2
      set send-version 1 2
    next
    edit "port2"
      set receive-version 1 2
      set send-version 1 2
    next
    edit "port3"
      set receive-version 1 2
      set send-version 1 2
    end
    edit "port4"
      set receive-version 1 2
      set send-version 1 2
    end
  config network
    edit 1
      set prefix 10.11.0.0 255.255.0.0
    next
    edit 2
      set prefix 10.12.0.0 255.255.0.0
    next
    edit 3
      set prefix 10.14.0.0 255.255.0.0
    next
    edit 4
      set prefix 172.20.120.0 255.255.255.0
    end
  config redistribute "static"
    set status enable
  end
end
```

Configuring other networking devices

In this example there are two groups of other devices on the the network — internal devices, and the ISP.

The first is the internal network devices on the Sales, R&D, and Accounting networks. This includes simple static routers, computers, printers and other network devices. Once the FortiGate units are configured, the internal static routers need to be configured using the internal network IP addresses. Otherwise there should be no configuration required.

The second group of devices is the ISP. This consists of the RIP router the FortiGate routers 2 and 3 connect to. You need to contact your ISP and ensure they have your information for your

network such as the IP addresses of the connecting RIP routers, what version of RIP your network supports, and what authentication (if any) is used.

Testing network configuration

Once the network has been configured, you need to test that it works as expected.

The two series of tests you need to run are to test the internal networks can communicate with each other, and that the internal networks can reach the internet.

Use ping, traceroute, and other networking tools to run these tests.

If you encounter problems, for troubleshooting help consult [“Troubleshooting RIP” on page 70](#).

RIPng – RIP and IPv6

RIP next generation, or RIPng, is the version of RIP that supports IPv6.

This is an example of a typical small network configuration using RIPng routing.

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate units running the RIPng dynamic routing protocol. Both FortiGate units are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network enabling it to reach the internet at all times.

This section includes the following topics:

- [Network layout and assumptions](#)
- [Configuring the FortiGate units system information](#)
- [Configuring RIPng on FortiGate units](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

Network layout and assumptions

Basic network layout

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate units running the RIPng dynamic routing protocol. Both FortiGate units are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network enabling it to reach the internet at all times.

All internal computers use RIP routing, so no static routing is required. And all internal computers use IPv6 addresses.


Where possible in this example, the default values will be used or the most general settings. This is intended to provide an easier configuration that will require less troubleshooting.

In this example the routers, networks, interfaces used, and IP addresses are as follows.

Table 10:Rip example network topology

Network	Router	Interface & Alias	IPv6 address
R&D	Router1	port1 (internal)	2002:A0B:6565:0:0:0:0:0
		port2 (ISP)	2002:AC14:7865:0:0:0:0:0
	Router2	port1 (internal)	2002:A0B:6566:0:0:0:0:0
		port2 (ISP)	2002:AC14:7866:0:0:0:0:0

Figure 5: Network topology for the IPV6 RIPng example



Assumptions

The following assumptions have been made concerning this example.

- All FortiGate units have 5.0 firmware, and are running factory default settings.
- All CLI and web-based manager navigation assumes the unit is running in NAT/Route operating mode, with VDOMs disabled.
- All FortiGate units have interfaces labelled port1 and port2 as required.
- All firewalls have been configured for each FortiGate unit to allow the required traffic to flow across interfaces.
- All network devices are support IPv6 and are running RIPng.

Configuring the FortiGate units system information

Each FortiGate unit needs IPv6 enabled, a new hostname, and interfaces configured.

To configure system information on Router1 - web-based manager

1. Go to *System > Dashboard > Status*.
2. For *Host name*, select *Change*.
3. Enter "Router1".
4. Go to *System > Admin > Settings*.
5. In *Display Options on GUI*, enable *IPv6*, and select *Apply*.
6. Go to *System > Network > Interfaces*.
7. Edit port1 (internal) interface.
8. Set the following information, and select *OK*.

Alias	internal
IP/Netmask	2002:A0B:6565::/0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

9. Edit port2 (ISP) interface.
10. Set the following information, and select *OK*.

Alias	ISP
IP/Netmask	2002:AC14:7865::/0
Administrative Access	HTTPS SSH PING
Description	ISP and internet
Administrative Status	Up

To configure system information on Router1 - CLI

```
config system global
  set hostname Router1
  set gui-ipv6 enable
end
config system interface
  edit port1
    set alias internal
    set allowaccess https ping ssh
    set description "Internal RnD network"
    config ipv6
      set ip6-address 2002:a0b:6565::/0
    end
  end
next
  edit port2
    set alias ISP
    set allowaccess https ping ssh
    set description "ISP and internet"
    config ipv6
      set ip6-address 2002:AC14:7865::
    end
  end
end
```

To configure system information on Router2 - web-based manager

1. Go to *System > Dashboard > Status*.
2. For *Host name*, select *Change*.
3. Enter "Router2".
4. Go to *System > Admin > Settings*.
5. In *Display Options on GUI*, enable *IPv6*, and select *Apply*.
6. Go to *System > Network > Interfaces*.
7. Edit port1 (internal) interface.
8. Set the following information, and select *OK*.

Alias	internal
IP/Netmask	2002:A0B:6566::/0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

9. Edit port2 (ISP) interface.
10. Set the following information, and select *OK*.

Alias	ISP
IP/Netmask	2002:AC14:7866::/0

Administrative Access	HTTPS SSH PING
Description	ISP and internet
Administrative Status	Up

To configure system information on Router2 - CLI

```

config system global
    set hostname Router2
    set gui-ipv6 enable
end
config system interface
    edit port1
        set alias internal
        set allowaccess https ping ssh
        set description "Internal RnD network"
        config ipv6
            set ip6-address 2002:a0b:6566::/0
        end
    next
    edit port2
        set alias ISP
        set allowaccess https ping ssh
        set description "ISP and internet"
        config ipv6
            set ip6-address 2002:AC14:7866::
        end
    end
end

```

Configuring RIPng on FortiGate units

Now that the interfaces are configured, you can configure RIPng on the FortiGate units.

There are only two networks and two interfaces to include — the internal network, and the ISP network. There is no redistribution, and no authentication. In RIPng there is no specific command to include a subnet in the RIP broadcasts. There is also no information required for the interfaces beyond including their name.

As this is a CLI only configuration, configure the ISP router and the other FortiGate unit as neighbors. This was not part of the previous example as this feature is not offered in the web-based manager. Declaring neighbors in the configuration like this will reduce the discovery traffic when the routers start up.

Since RIPng is not supported in the web-based manager, this section will only be entered in the CLI.

To configure RIPng on Router1 - CLI

```
config router ripng
  config interface
    edit port1
    next
    edit port2
    end
  config neighbor
    edit 1
      set interface port1
      set ipv6 2002:a0b:6566::/0
    next
    edit 2
      set interface port2
      set ipv6 2002:AC14:7805::/0
    end
```

To configure RIPng on Router2 - CLI

```
config router ripng
  config interface
    edit port1
    next
    edit port2
    end
  config neighbor
    edit 1
      set interface port1
      set ipv6 2002:a0b:6565::/0
    next
    edit 2
      set interface port2
      set ipv6 2002:AC14:7805::/0
    end
```

Configuring other network devices

The other devices on the internal network all support IPv6, and are running RIPng where applicable. They only need to know the internal interface network addresses of the FortiGate units.

The ISP routers need to know the FortiGate unit information such as IPv6 addresses.

Testing the configuration

In addition to normal testing of your network configuration, you must also test the IPv6 part of this example.

For troubleshooting problems with your network, see the FortiOS Handbook Troubleshooting chapter.

For troubleshooting problems with RIP, see [“Troubleshooting RIP”](#) on page 70.

Testing the IPv6 RIPng information

There are some commands to use when checking that your RIPng information is correct on your network. These are useful to check on your RIPng FortiGate units on your network. Comparing the output between devices will help you understand your network better, and also track down any problems.

```
diagnose ipv6 address list
```

View the local scope IPv6 addresses used as next-hops by RIPng on the FortiGate unit.

```
diagnose ipv6 route list
```

View ipv6 addresses that are installed in the routing table.

```
get router info6 routing-table
```

View the routing table. This information is almost the same as the previous command (diagnose ipv6 route list) however it is presented in an easier to read format.

```
get router info6 rip interface external
```

View brief output on the RIP information for the interface listed. The information includes if the interface is up or down, what routing protocol is being used, and whether passive interface or split horizon are enabled.

```
get router info6 neighbor-cache list
```

View the IPv6/MAC address mapping. This also displays the interface index and name associated with the address.

Border Gateway Protocol (BGP)

This section describes Border Gateway Protocol (BGP).

The following topics are included in this section:

- [BGP background and concepts](#)
- [Troubleshooting BGP](#)
- [Dual-homed BGP example](#)
- [Redistributing and blocking routes in BGP](#)

BGP background and concepts

The border gateway protocol contains two distinct subsets — internal BGP (iBGP) and external BGP (eBGP). iBGP is intended for use within your own networks. eBGP is used to connect many different networks together, and is the main routing protocol for the Internet backbone. FortiGate units support iBGP, and eBGP only for communities.

The following topics are included in this section:

- [Background](#)
- [Parts and terminology of BGP](#)
- [How BGP works](#)

Background

BGP was first used in 1989. The current version, BGP-4, was released in 1995 and is defined in RFC 1771. That RFC has since been replaced by the more recent RFC 4271. The main benefits of BGP-4 are classless inter-domain routing, and aggregate routes. BGP is the only routing protocol to use TCP for a transport protocol. Other routing protocols use UDP.

BGP makes routing decisions based on path, network policies and rulesets instead of the hop-count metric as RIP does, or cost-factor metrics as OSPF does.

BGP-4+ supports IPv6. It was introduced in RFC 2858 and RFC 2545.

BGP is the routing protocol used on the Internet. It was designed to replace the old Exterior Gateway Protocol (EGP) which had been around since 1982, and was very limited. In doing so, BGP enabled more networks to take part in the Internet backbone to effectively decentralize it and make the Internet more robust, and less dependent on a single ISP or backbone network.

Parts and terminology of BGP

In a BGP network, there are some terms that need to be explained before going ahead. Some parts of BGP are not explained here as they are common to other dynamic routing protocols as well. When determining your network topology, note that the number of available or supported routes is not set by the configuration but depends on your FortiGate's available memory. For more information on parts of BGP that are not listed here, see [“Dynamic routing terminology” on page 53](#).

BGP and IPv6

FortiGate units support IPv6 over BGP using the same `config router bgp` command as IPv4, but different subcommands.

The main CLI keywords have IPv6 equivalents that are identified by the “6” on the end of the keyword, such as with `config network6` or `set allowas-in6`. For more information about IPv6 BGP keywords, see the [FortiGate CLI Reference](#).

IPv6 BGP commands include:

```
config router bgp
  set activate6 {enable | disable}
  set allowas-in6 <max_num_AS_integer>
  set allowas-in-enable6 {enable | disable}
  set as-override6 {enable | disable}
  set attribute-unchanged6 [as-path] [med] [next-hop]
  set capability-default-originate6 {enable | disable}
  set capability-graceful-restart6 {enable | disable}
  set capability-orf6 {both | none | receive | send}
  set default-originate-route-map6 <routemap_str>
  set distribute-list-in6 <access-list-name_str>
  set distribute-list-out6 <access-list-name_str>
  set filter-list-in6 <aspath-list-name_str>
  set filter-list-out6 <aspath-list-name_str>
  set maximum-prefix6 <prefix_integer>
  set maximum-prefix-threshold6 <percentage_integer>
  set maximum-prefix-warning-only6 {enable | disable}
  set next-hop-self6 {enable | disable}
  set prefix-list-in6 <prefix-list-name_str>
  set prefix-list-out6 <prefix-list-name_str>
  set remove-private-as6 {enable | disable}
  set route-map-in6 <routemap-name_str>
  set route-map-out6 <routemap-name_str>
  set route-reflector-client6 {enable | disable}
  set route-server-client6 {enable | disable}
  set send-community6 {both | disable | extended | standard}
  set soft-reconfiguration6 {enable | disable}
  set unsuppress-map6 <route-map-name_str>
  config network6
  config redistribute6
end
```

Roles of routers in BGP networks

Dynamic routing has a number of different roles routers can fill such as those covered in [“Dynamic routing terminology”](#) on page 53. BGP has a number of custom roles that routers can fill. These include:

- Speaker routers
- Peer routers or neighbors
- Route reflectors (RR)

Speaker routers

Any router configured for BGP is considered a BGP speaker. This means that a speaker router advertises BGP routes to its peers.

Any routers on the network that are not speaker routers, are not treated as BGP routers.

Peer routers or neighbors

In a BGP network, all neighboring BGP routers or peer routers are routers that are connected to your FortiGate unit. Your FortiGate unit learns about all other routers through these peers.

You need to manually configure BGP peers on your FortiGate unit as neighbors. Otherwise these routers will not be seen as peers, but instead as simply other routers on the network that don't support BGP. You can optionally use MD5 authentication to password protect BGP sessions with those neighbors. (see RFC 2385).

You can configure up to 1000 BGP neighbors on your FortiGate unit. You can clear all or some BGP neighbor connections (sessions) using the `execute router clear bgp` command.

For example, if you have 10 routes in the BGP routing table and you want to clear the specific route to IP address 10.10.10.1, enter the command:

```
execute router clear bgp ip 10.10.10.1
```

To remove all routes for AS number 650001, enter the command:


```
execute router clear bgp as 650001
```

To remove route flap dampening information for the 10.10.0.0/16 subnet, enter the command:

```
execute router clear bgp dampening 10.10.0.0/16
```

In Figure 1, Router A is directly connected to five other routers in a network that contains 12 routers overall. These routers, the ones in the blue circle, are Router A's peers or neighbors.

Figure 6: Router A and its 5 peer routers



Router A's peer routers

As a minimum, when configuring BGP neighbors you must enter their IP address, and the AS number (*remote-as*). This is all the information the web-based manager interface allows you to enter for a neighbor.

The BGP commands related to neighbors are quite extensive and include:

```
config router bgp
  config neighbor
    edit <neighbor_address_ipv4>
      set activate {enable | disable}
      set advertisement-interval <seconds_integer>
      set allowas-in <max_num_AS_integer>
      set allowas-in-enable {enable | disable}
      set as-override {enable | disable}
      set attribute-unchanged [as-path] [med] [next-hop]
      set bfd {enable | disable}
      set capability-default-originate {enable | disable}
      set capability-dynamic {enable | disable}
      set capability-graceful-restart {enable | disable}
      set capability-orf {both | none | receive | send}
      set capability-route-refresh {enable | disable}
      set connect-timer <seconds_integer>
      set description <text_str>
      set distribute-list-in <access-list-name_str>
      set distribute-list-out <access-list-name_str>
      set dont-capability-negotiate {enable | disable}
      set ebgp-enforce-multihop {enable | disable}
      set ebgp-multihop {enable | disable}
      set ebgp-multihop-ttl <seconds_integer>
      set filter-list-in <aspath-list-name_str>
      set filter-list-out <aspath-list-name_str>
      set holdtime-timer <seconds_integer>
      set interface <interface-name_str>
      set keep-alive-timer <seconds_integer>
      set maximum-prefix <prefix_integer>
      set maximum-prefix-threshold <percentage_integer>
      set maximum-prefix-warning-only {enable | disable}
      set next-hop-self {enable | disable}
      set passive {enable | disable}
      set password <string>
      set prefix-list-in <prefix-list-name_str>
      set prefix-list-out <prefix-list-name_str>
      set remote-as <id_integer>
      set remove-private-as {enable | disable}
      set retain-stale-time <seconds_integer>
      set route-map-in <routemap-name_str>
      set route-map-out <routemap-name_str>
      set route-reflector-client {enable | disable}
      set route-server-client {enable | disable}
      set send-community {both | disable | extended | standard}
      set shutdown {enable | disable}
      set soft-reconfiguration {enable | disable}
      set strict-capability-match {enable | disable}
```

```
        set unsuppress-map <route-map-name_str>
        set update-source <interface-name_str>
        set weight <weight_integer>
    end
end
end
```

Route reflectors (RR)

Route reflectors in BGP concentrate route updates so other routers need only talk to the route reflectors to get all the updates. This results in smaller routing tables, fewer connections between routers, faster responses to network topology changes, and less administration bandwidth. BGP route reflectors are defined in RFC 1966.


In a BGP route reflector configuration, the AS is divided into different clusters that each include client and reflector routers. The client routers supply the reflector routers with the client's route updates. The reflectors pass this information along to other route reflectors and border routers. Only the reflectors need to be configured, not the clients — the clients will find the closest reflector and communicate with it automatically. The reflectors communicate with each other as peers. FortiGate units can be configured as either reflectors or clients.

Since route reflectors are processing more than the client routers, the reflectors should have more resources to handle the extra workload.

Smaller networks running BGP typically don't require route reflectors (RR). However, RR is a useful feature for large companies, where their AS may include 100 routers or more. For example, for a full mesh 20 router configuration within an AS, there would have to be 190 unique BGP sessions — just for routing updates within the AS. The number of sessions jumps to 435 sessions for just 30 routers, or 4950 sessions for 100 routers. From these numbers, it's plain that updating this many sessions will quickly consume the limited bandwidth and processing resources of the routers involved.

The following diagram illustrates how route reflectors can improve the situation when only six routers are involved. The AS without route reflectors requires 15 sessions between the routers. In the AS with route reflectors, the two route reflectors receive route updates from the reflector clients (unlabeled routers in the diagram) in their cluster as well as other route reflectors and pass them on to the border router. The RR configuration only requires six sessions. This example shows a reduction of 60% in the number of required sessions.

Figure 7: Required sessions within an AS with and without route reflectors



The BGP commands related to route reflectors includes:

```
config router bgp
  config neighbor
    set route-reflector-client {enable | disable}
    set route-server-client {enable | disable}
  end
end
```

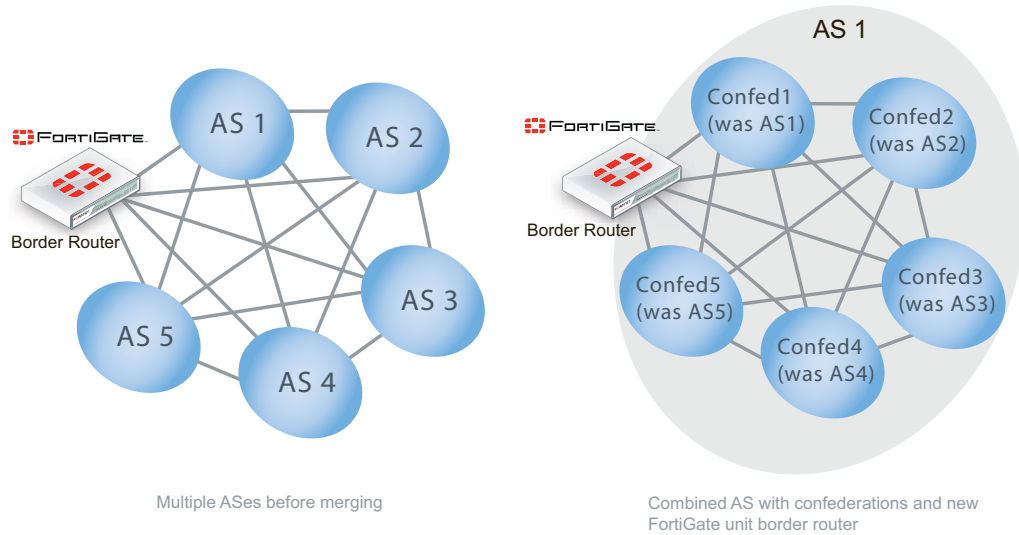
Confederations

Confederations were introduced to reduce the number of BGP advertisements on a segment of the network, and reduce the size of the routing tables. Confederations essentially break up an AS into smaller units. Confederations are defined in RFC 3065 and RFC 1965.

Within a confederation, all routers communicate with each other in a full mesh arrangement. Communications between confederations is more like inter-AS communications in that many of the attributes are changed as they would be for BGP communications leaving the AS, or eBGP.

Confederations are useful when merging ASs. Each AS being merged can easily become a confederation, requiring few changes. Any additional permanent changes can then be implemented over time as required. The figure below shows the group of ASs before merging, and the corresponding confederations afterward as part of the single AS with the addition of a new border router. It should be noted that after merging if the border router becomes a route reflector, then each confederation only needs to communicate with one other router, instead of five others.

Figure 8: AS merging using confederations



Confederations and route reflectors perform similar functions — they both sub-divide large ASes for more efficient operation. They differ in that route reflector clusters can include routers that are not members of a cluster, where routers in a confederation must belong to that confederation. Also, confederations place their confederation numbers in the AS_PATH attribute making it easier to trace.

It is important to note that while confederations essentially create sub-ASs, all the confederations within an AS appear as a single AS to external ASs.

Confederation related BGP commands include:

```
config router bgp
    set confederation-identifier <peerid_integer>
end
```

Network Layer Reachability Information (NLRI)

Network Layer Reachability Information (NLRI) is unique to BGP-4. It is sent as part of the update messages sent between BGP routers, and contains information necessary to supernet, or aggregate route, information. The NLRI includes the length and prefix that when combined are the address of the aggregated routes referred to.

There is only one NLRI entry per BGP update message.

BGP attributes

Each route in a BGP network has a set of attributes associated with it. These attributes define the route, and are modified as required along the route.

BGP can work well with mostly default settings, but if you are going to change settings you need to understand the roles of each attribute and how they affect those settings.

The BGP attributes include:

AS_PATH	A list of ASes a route has passed through. See “AS_PATH” on page 105.
MULTI_EXIT_DESC (MED)	Which router to use to exit an AS with more than one external connection. See “MULTI_EXIT_DESC (MED)” on page 106.

COMMUNITY	Used to apply attributes to a group of routes. See “COMMUNITY” on page 106.
NEXT_HOP	Where the IP packets should be forwarded to, like a gateway in static routing. See “NEXT_HOP” on page 106.
ATOMIC_AGGREGATE	Used when routes have been summarized to tell downstream routers not to de-aggregate the route. See “ATOMIC_AGGREGATE” on page 107.
ORIGIN	Used to determine if the route is from the local AS or not. See “ORIGIN” on page 107.
LOCAL_PREF	Used only within an AS to select the best route to a location (like MED)



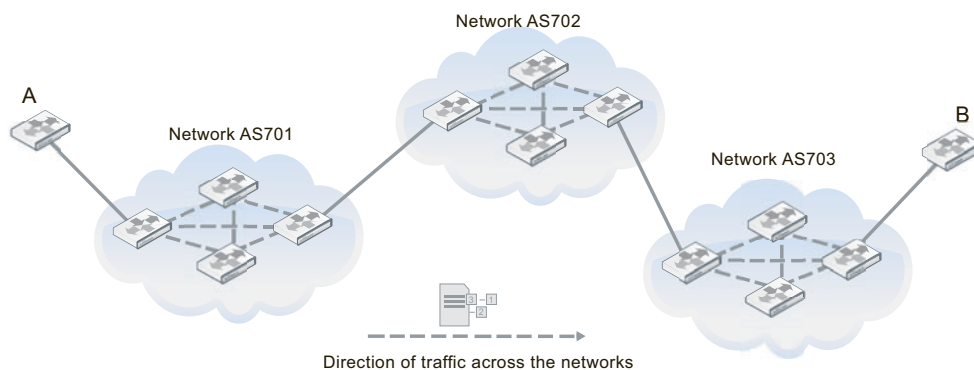
Inbound policies on FortiGate units can change the NEXT-HOP, LOCAL-PREF, MED and AS-PATH attributes of an internal BGP (iBGP) route for its local route selection purposes. However, outbound policies on the unit cannot affect these attributes.

AS_PATH

AS_PATH is the BGP attribute that keeps track of each AS a route advertisement has passed through. AS_PATH is used by confederations and by exterior BGP (EBGP) to help prevent routing loops. A router knows there is a loop if it receives an AS_PATH with that routers AS in it. The figure below shows the route between router A and router B. The AS_PATH from A to B would read 701,702,703 for each AS the route passes through.

As of the start of 2010, the industry upgraded from 2-byte to 4-byte AS_PATHs. This upgrade was due to the imminent exhaustion of 2-byte AS_PATH numbers. FortiOS supports 4-byte AS_PATHs in its BGP implementation.

Figure 9: AS_PATH of 701,702, 703 between routers A and B



The BGP commands related to AS_PATH include:

```
config router bgp
  set bestpath-as-path-ignore {enable | disable}
end
```

MULTI_EXIT_DESC (MED)

BGP AS systems can have one or more routers that connect them to other ASes. For ASes with more than one connecting router, the Multi-Exit Discriminator (MED) lists which router is best to use when leaving the AS. The MED is based on attributes such as delay. It is a recommendation only, as some networks may have different priorities.

BGP updates advertise the best path to a destination network. When the FortiGate unit receives a BGP update, the FortiGate unit examines the Multi-Exit Discriminator (MED) attribute of potential routes to determine the best path to a destination network before recording the path in the local FortiGate unit routing table.

FortiGate units have the option to treat any routes without an MED attribute as the worst possible routing choice. This can be useful because a lack of MED information is a lack of routing information which can be suspicious — possibly a hacking attempt or an attack on the network. At best it signifies an unreliable route to select.

The BGP commands related to MED include:

```
config router bgp
    set always-compare-med {enable | disable}
    set bestpath-med-confed {enable | disable}
    set bestpath-med-missing-as-worst {enable | disable}
    set deterministic-med {enable | disable}
    config neighbor
        set attribute-unchanged [as-path] [med] [next-hop]
    end
end
```

COMMUNITY

A community is a group of routes that have the same routing policies applied to them. This saves time and resources. A community is defined by the COMMUNITY attribute of a BGP route.

The FortiGate unit can set the COMMUNITY attribute of a route to assign the route to predefined paths (see RFC 1997). The FortiGate unit can examine the COMMUNITY attribute of learned routes to perform local filtering and/or redistribution.

The BGP commands related to COMMUNITY include:

```
config router bgp
    set send-community {both | disable | extended | standard}
end
```

NEXT_HOP

The NEXT_HOP attribute says what IP address the packets should be forwarded to next. Each time the route is advertised, this value is updated. The NEXT_HOP attribute is much like a gateway in static routing.

FortiGate units allow you to change the advertising of the FortiGate unit's IP address (instead of the neighbor's IP address) in the NEXT_HOP information that is sent to IBGP peers. This is changed with the config neighbor, set next-hop-self command.

The BGP commands related to NEXT_HOP include:

```
config router bgp
  config neighbor
    set attribute-unchanged [as-path] [med] [next-hop]
    set next-hop-self {enable | disable}
  end
end
```

ATOMIC_AGGREGATE

The ATOMIC_AGGREGATE attribute is used when routes have been summarized. It indicates which AS and which router summarize the routes. It also tells downstream routers not to de-aggregate the route. Summarized routes are routes with similar information that have been combined, or aggregated, into one route that is easier to send in updates. When it reaches its destination, the summarized routes are split back up into the individual routes.

Your FortiGate unit doesn't specifically set this attribute in the BGP router command, but it is used in the route map command.

The commands related to ATOMIC_AGGREGATE include:

```
config router route-map
  edit <route_map_name>
    config rule
      edit <route_map_rule_id>
        set set-aggregator-as <id_integer>
        set set-aggregator-ip <address_ipv4>
        set set-atomic-aggregate {enable | disable}
      end
    end
  end
end
```

ORIGIN

The ORIGIN attribute records where the route came from. The options can be IBGP, EBGP, or incomplete. This information is important because internal routes (IBGP) are by default higher priority than external routes (EBGP). However incomplete ORIGINS are the lowest priority of the three.

The commands related to ORIGIN include:

```
config router route-map
  edit <route_map_name>
    set comments <string>
    config rule
      edit <route_map_rule_id>
        set match-origin {egp | igp | incomplete | none}
      end
    end
  end
end
```

How BGP works

BGP is a link-state routing protocol and keeps link-state information about the status of each network link it has connected. A BGP router receives information from its peer routers that have

been defined as neighbors. BGP routers listen for updates from these configured neighboring routers on TCP port 179.

A BGP router is a finite state machine with six various states for each connection. As two BGP routers discover each other, and establish a connection they go from the idle state, through the various states until they reach the established state. An error can cause the connection to be dropped and the state of the router to be reset to either active or idle. These errors can be caused by: TCP port 179 not being open, a random TCP port above port 1023 not being open, the peer address being incorrect, or the AS number being incorrect.

When BGP routers start a connection, they negotiate which (if any) optional features will be used such as multiprotocol extensions that can include IPv6 and VPNs.

IBGP versus EBGp

When you read about BGP, often you see EBGp or IBGP mentioned. These are both BGP routing, but BGP used in different roles. Exterior BGP (EBGP) involves packets crossing multiple autonomous systems (ASes) where interior BGP (IBGP) involves packets that stay within a single AS. For example the AS_PATH attribute is only useful for EBGp where routes pass through multiple ASes.

These two modes are important because some features of BGP are only used for one of EBGp or IBGP. For example confederations are used in EBGp, and route reflectors are only used in IBGP. Also routes learned from IBGP have priority over EBGp learned routes.

FortiGate units have some commands specific to EBGp. These include:

- automatically resetting the session information to external peers if the connection goes down — `set fast-external-failover {enable | disable}`
- setting an administrative distance for all routes learned from external peers (must also configure local and internal distances if this is set) — `set distance-external <distance_integer>`
- enforcing EBGp multihops and their TTL (number of hops) — `set ebgp-enforce-multihop {enable | disable}` and `set ebgp-multihop-ttl <seconds_integer>`

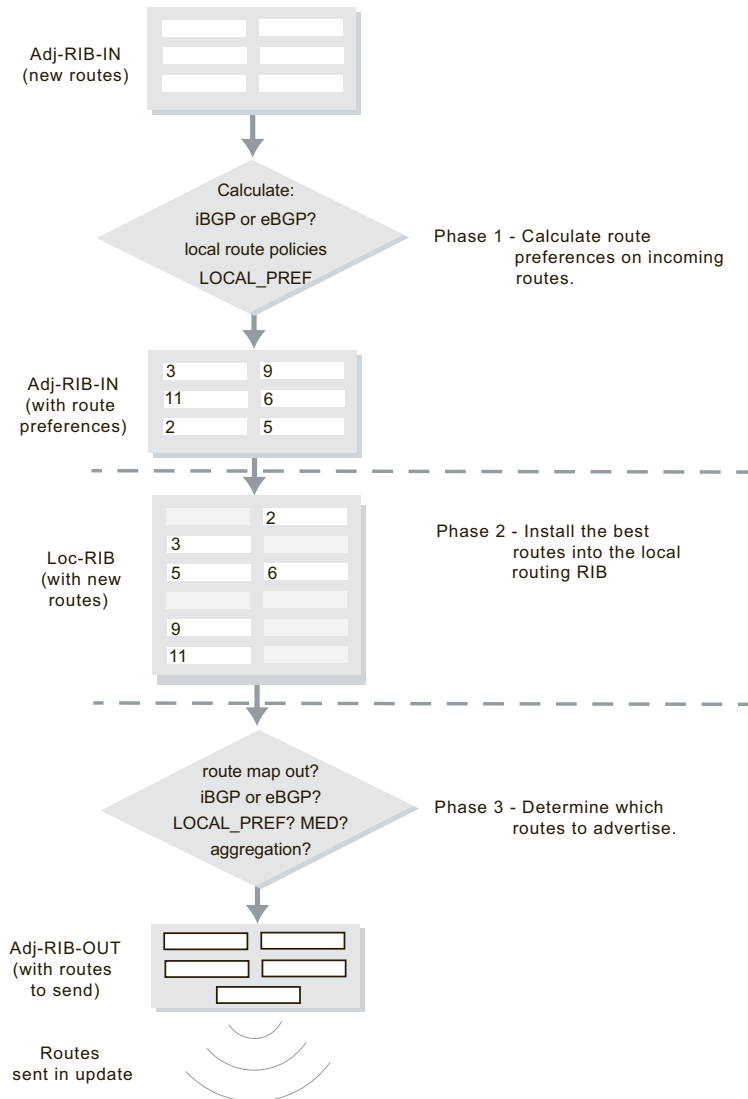
BGP path determination — which route to use

Firstly, recall that the number of available or supported routes is not set by the configuration but depends on your FortiGate's available memory. All learned routes and their attributes come into the BGP router in raw form. Before routes are installed in the routing table or are advertised to other routers, three levels of decisions must be made.

The three phases of BGP best path determination do not change. However, some manufacturers have added more information to the process, such as Cisco's WEIGHT attribute to enable an administrator to force one route's selection over another.

There is one Adj-RIB-IN and Adj-RIB-OUT for each configured neighbor. They are updated when the FortiGate unit receives BGP updates, or when the FortiGate unit sends out BGP updates.

Figure 10: Three phases of BGP routing decision



Decision phase 1

At this phase, the decision is to calculate how preferred each route and its NRI are the Adjacent Routing Information Base Incoming (Adj-RIBs-In) compared to the other routes. For internal routes (IBGP), policy information or LOCAL_PREF is used. For external peer learned routes, it is based strictly on policy. These rules set up a list of which routes are most preferred going into Phase 2.

Decision phase 2

Phase 2 involves installing the best route to each destination into the local Routing Information Base (Loc-RIB). Effectively, the Loc-RIB is the master routing table. Each route from Phase 1 has their NEXT_HOP checked to ensure the destination is reachable. If it is reachable, the

AS_PATH is checked for loops. After that, routes are installed based on the following decision process:

- If there is only one route to a location, it is installed.
- If multiple routes to the same location, use the most preferred route from Level 1.
- If there is a tie, break the tie based on the following in descending order of importance: shortest AS_PATH, smallest ORIGIN number, smallest MED, EBGp over IBGP, smallest metric or cost for reaching the NEXT_HOP, BGP identifier, and lowest IP address.

Note that the new routes that are installed into the Loc-RIB are in addition to any existing routes in the table. Once Phase 2 is completed the Loc-RIB will consist of the best of both the new and older routes.

Decision phase 3

Phase 3 is route distribution or dissemination. This is the process of deciding which routes the router will advertise. If there is any route aggregation or summarizing, it happens here. Also any route filtering from route maps happens here.

Once Phase 3 is complete, an update can be sent out to update the neighbor of new routes.

Aggregate routes and addresses

BGP4 allows classless routing, which uses netmasks as well as IP addresses. This classless routing enables the configuration of aggregate routes by stating the address bits the aggregated addresses have in common. For more information, see [“Aggregated routes and addresses” on page 53](#).

The ATOMIC_AGGREGATE attribute informs routers that the route has been aggregated, and should not be de-aggregated. An associated AGGREGATOR attribute include the information about the router that did the aggregating including its AS.

The BGP commands associated with aggregate routes and addresses are:

```
config router bgp
  config aggregate-address
    edit <aggr_addr_id>
      set as-set {enable | disable}
      set prefix <address_ipv4mask>
      set summary-only {enable | disable}
    end
  config aggregate-address6
    edit <aggr_addr_id>
      set as-set {enable | disable}
      set prefix6 <address_ipv6mask>
      set summary-only {enable | disable}
    end
  end
end
```

Troubleshooting BGP

There are some features in BGP that are used to deal with problems that may arise. Typically the problems with a BGP network that has been configured, involve routes going offline frequently. This is called route flap and causes problems for the routers using that route.

Clearing routing table entries

To see if a new route is being properly added to the routing table, you can clear all or some BGP neighbor connections (sessions) using the `execute router clear bgp` command.

For example, if you have 10 routes in the BGP routing table and you want to clear the specific route to IP address 10.10.10.1, enter the command:

```
execute router clear bgp ip 10.10.10.1
```

To remove all routes for AS number 650001, enter the command:

```
execute router clear bgp as 650001
```

Route flap

When routers or hardware along a route go offline and back online that is called a route flap. Flapping is the term if these outages continue, especially if they occur frequently.

Route flap is a problem in BGP because each time a peer or a route goes down, all the peer routers that are connected to that out-of-service router advertise the change in their routing tables which creates a lot of administration traffic on the network. And the same traffic happens again when that router comes back online. If the problem is something like a faulty network cable that wobbles on and offline every 10 seconds, there could easily be overwhelming amounts of routing updates sent out unnecessarily.

Another possible reason for route flap occurs with multiple FortiGate units in HA mode. When an HA cluster fails over to the secondary unit, other routers on the network may see the HA cluster as being offline resulting in route flap. While this doesn't occur often, or more than once at a time, it can still result in an interruption in traffic which is unpleasant for network users. The easy solution for this problem is to increase the timers on the HA cluster, such as TTL timers, so they do not expire during the failover process. Also configuring graceful restart on the HA cluster will help with a smooth failover.

The first method of dealing with route flap should be to check your hardware. If a cable is loose or bad, it can easily be replaced and eliminate the problem. If an interface on the router is bad, either avoid using that interface or swap in a functioning router. If the power source is bad on a router, either replace the power supply or use a power conditioning backup power supply. These quick and easy fixes can save you from configuring more complex BGP options. However if the route flap is from another source, configuring BGP to deal with the outages will ensure your network users uninterrupted service.

Some methods of dealing with route flap in BGP include:

- [Holddown timer](#)
- [Dampening](#)
- [Graceful restart](#)
- [Bi-directional forwarding detection \(BFD\)](#)

Holddown timer

The first line of defence to a flapping route is the hold down timer. This timer reduces how frequently a route going down will cause a routing update to be broadcast.

Once activated, the holddown timer won't allow the FortiGate unit to accept any changes to that route for the duration of the timer. If the route flaps five times during the timer period, only the first outage will be recognized by the FortiGate unit — for the duration of the other outages there will be no changes because the Fortigate unit is essentially treating this router as down. After the timer expires, if the route is still flapping it will happen all over again.

Even if the route isn't flapping — if it goes down, comes up, and stays back up — the timer still counts down and the route is ignored for the duration of the timer. In this situation the route will be seen as down longer than it really is, but there will be only the one set of route updates. This is not a problem in normal operation because updates are not frequent.

Also the potential for a route to be treated as down when it is really up can be viewed as a robustness feature. Typically you do not want most of your traffic being routed over an unreliable route. So if there is route flap going on, it is best to avoid that route if you can. This is enforced by the holddown timer.

How to configure the holddown timer

There are three different route flapping situations that can occur: the route goes up and down frequently, the route goes down and back up once over a long period of time, or the route goes down and stays down for a long period of time. These can all be handled using the holddown timer.

For example, your network has two routes that you want to set the holddown timer for. One is your main route (to 10.12.101.4) that all your Internet traffic goes through, and it can't be down for long if it's down. The second is a low speed connection to a custom network that is used infrequently (to 10.13.101.4). The holddown timer for the main route should be fairly short, let's say 60 seconds instead of the default 180 seconds. The second route timer can be left at the default or even longer since it is rarely used. In your BGP configuration this looks like:

```
config router bgp
  config neighbor
  edit 10.12.101.4
    set holddown-timer 60
  next
  edit 10.13.101.4
    set holddown-timer 180
  next
end
end
```

Dampening

Dampening is a method used to limit the amount of network problems due to flapping routes. With dampening the flapping still occurs, but the peer routers pay less and less attention to that route as it flaps more often. One flap doesn't start dampening, but the second starts a timer where the router will not use that route — it is considered unstable. If the route flaps again before the timer expires, the timer continues to increase. There is a period of time called the reachability half-life after which a route flap will only be suppressed for half the time. This half-life comes into effect when a route has been stable for a while but not long enough to clear all the dampening completely. For the flapping route to be included in the routing table again, the suppression time must expire.

If the route flapping was temporary, you can clear the flapping or dampening from the FortiGate units cache by using one of the `execute router clear bgp` commands:

```
execute router clear bgp dampening {<ip_address> | <ip/netmask>}
```

or

```
execute router clear bgp flap-statistics {<ip> | <ip/netmask>}
```

For example, to remove route flap dampening information for the 10.10.0.0/16 subnet, enter the command:

```
execute router clear bgp dampening 10.10.0.0/16
```


The BGP commands related to route dampening are:

```
config router bgp
  set dampening {enable | disable}
  set dampening-max-suppress-time <minutes_integer>
  set dampening-reachability-half-life <minutes_integer>
  set dampening-reuse <reuse_integer>
  set dampening-route-map <routemap-name_str>
  set dampening-suppress <limit_integer>
  set dampening-unreachability-half-life <minutes_integer>
end
```

Graceful restart

BGP4 has the capability to gracefully restart.

In some situations, route flap is caused by routers that appear to be offline but the hardware portion of the router (control plane) can continue to function normally. One example of this is when some software is restarting or being upgraded, but the hardware can still function normally.

Graceful restart is best used for these situations where routing will not be interrupted, but the router is unresponsive to routing update advertisements. Graceful restart does not have to be supported by all routers in a network, but the network will benefit when more routers support it.



FortiGate HA clusters can benefit from graceful restart. When a failover takes place, the HA cluster will advertise it is going offline, and will not appear as a route flap. It will also enable the new HA main unit to come online with an updated and usable routing table — if there is a flap the HA cluster routing table will be out of date.

For example, your FortiGate unit is one of four BGP routers that send updates to each other. Any of those routers may support graceful starting—when a router plans to go offline, it will send out a message to its neighbors how long it expects to be before being back online. That way its neighbor routers don't remove it from their routing tables. However if that router isn't back online when expected, the routers will mark it offline. This prevents routing flap and its associated problems.

Scheduled time offline

Graceful restart is a means for a router to advertise it is going to have a scheduled shutdown for a very short period of time. When neighboring routers receive this notice, they will not remove that router from their routing table until after a set time elapses. During that time if the router comes back online, everything continues to function as normal. If that router remains offline longer than expected, then the neighboring routers will update their routing tables as they assume that router will be offline for a long time.

FortiGate units support both graceful restart of their own BGP routing software, and also neighboring BGP routers.

For example, if a neighbor of your FortiGate unit, with an IP address of 172.20.120.120, supports graceful restart, enter the command:

```
config router bgp
  config neighbor
  edit 172.20.120.120
    set capability-graceful-restart enable
  end
end
```

If you want to configure graceful restart on your FortiGate unit where you expect the Fortigate unit to be offline for no more than 2 minutes, and after 3 minutes the BGP network should consider the FortiGate unit offline, enter the command:

```
config router bgp
  set graceful-restart enable
  set graceful-restart-time 120
  set graceful-stalepath-time 180
end
```

The BGP commands related to BGP graceful restart are:

```
config router bgp
  set graceful-restart { disable| enable}
  set graceful-restart-time <seconds_integer>
  set graceful-stalepath-time <seconds_integer>
  set graceful-update-delay <seconds_integer>
  config neighbor
    set capability-graceful-restart {enable | disable}
  end
end

execute router restart
```

Before the restart, the router sends its peers a message to say it is restarting. The peers mark all the restarting router's routes as stale, but they continue to use the routes. The peers assume the router will restart and check its routes and take care of them if needed after the restart is complete. The peers also know what services the restarting router can maintain during its restart. After the router completes the restart, the router sends its peers a message to say it is done restarting.

Bi-directional forwarding detection (BFD)

Bi-directional Forwarding Detection (BFD) is a protocol used to quickly locate hardware failures in the network. Routers running BFD communicate with each other, and if a timer runs out on a connection then that router is declared down. BFD then communicates this information to the routing protocol and the routing information is updated.

While BGP can detect route failures, BFD can be configured to detect these failures more quickly allowing faster responses and improved convergence. This can be balanced with the bandwidth BFD uses in its frequent route checking.

Configurable granularity

BFD can run on the entire FortiGate unit, selected interfaces, or on BGP for all configured interfaces. The hierarchy allows each lower level to override the upper level's BFD setting. For example, if BFD was enabled for the FortiGate unit, it could be disabled only for a single

interface or for BGP. For information about FortiGate-wide BFD options, see config system settings in the *FortiGate CLI Reference*.

BFD can only be configured through the CLI.

The BGP commands related to BFD are:

```
config system {setting | interface}
    set bfd {enable | disable | global}
    set bfd-desired-mix-tx <milliseconds>
    set bfd-detect-mult <multiplier>
    set bfd-required-mix-rx <milliseconds>
    set bfd-dont-enforce-src-port {enable | disable}

config router bgp
    config neighbor
        edit <neighbor_address_ipv4>
            set bfd {enable | disable}
        end
    end

get router info bfd neighbor
execute router clear bfd session <src_ipv4> <dst_ipv4> <interface>
```

The `config system` commands allow you to configure whether BFD is enabled in a particular unit/vdom or individual interface, and how often the interface requires sending and receiving of BFD information.

The `config router bgp` commands allow you to set the addresses of the neighbor units that are also running BFD. Both units must be configured with BFD in order to make use of it.

Dual-homed BGP example

This is an example of a small network that uses BGP routing connections to two ISPs. This is a common configuration for companies that need redundant connections to the Internet for their business.

This configuration is for a small company connected to two ISPs. The company has one main office, the Head Office, and uses static routing for internal routing on that network.

Both ISPs use BGP routing, and connect to the Internet directly. They want the company to connect to the ISP networks using BGP. They also use graceful restart to prevent unneeded updates, and use smaller timer values to detect network failures faster.

As can be expected, the company wants to keep their BGP configuration relatively simple and easy to manage. The current configuration has only 3 routers to worry about — the 2 ISP border routers, and the FortiGate unit. This means the FortiGate unit will only have two neighbour routers to configure.

This configuration has the added benefit of being easy to expand if the Company wants to add a remote office in the future.

To keep the configuration simple, the Company is allowing only HTTP, HTTPS, FTP, and DNS traffic out of the local network. This will allow employees access to the Internet and their web-mail.

This section includes the following topics:

- [Network layout and assumptions](#)
- [Configuring the FortiGate unit](#)
- [Configuring other networking devices](#)
- [Testing this configuration](#)

Why dual home?

Dual homing means having two separate independent connections to the Internet. Servers in this configuration have also been called bastion hosts and can include DNS servers which require multiple connections.

Benefits of dual homing can include:

- redundant Internet connection that essentially never fails
- faster connections through one ISP or the other for some destinations, such as other clients of those ISPs
- load balancing traffic to your Company network
- easier to enable more traffic through two connections than upgrading one connection to bigger bandwidth
- easier to create protection policies for different traffic through a specific ISP

Some companies require reliable internet access at all times as part of their business. Consider a doctor operating remotely who has their Internet connection fail — the consequences could easily be life or death.

Dual homing is extra expense for the second ISP connection, and more work to configure and maintain the more complex network topology.

Potential dual homing issues

BGP comes with load balancing issues, and dual homing is the same category. BGP does not inherently deal well with load balancing, or getting default routes through BGP. Ideally one connect may be best for certain destinations, but it may not have that traffic routed to it making the load balancing less than perfect. This kind of fine tuning can be very time consuming, and usually results in a best effort situation.

When dual homing is not configured properly, your network may become a link between your ISPs and result in very high traffic between the ISPs that does not originate from your network. The problems with this situation are that your traffic may not have the bandwidth it needs, and you will be paying for a large volume of traffic that is not yours. This problem can be solved by not broadcasting or redistributing BGP routes between the ISPs.

If you learn your default routes from the ISPs in this example, you may run into an asymmetric routing problem where your traffic loops out one ISP and back to you through the other ISP. If you think this may be happening you can turn on asymmetric routing on the FortiGate unit (config system settings, set asymmetric enable) to verify that really is the problem. Turn this feature off once this is established since it disables many features on the FortiGate by disabling stateful inspection. Solutions for this problem can include using static routes for default routes instead of learning them through BGP, or configuring VDOMs on your FortiGate unit to provide a slightly different path back that is not a true loop.


Network layout and assumptions

The network layout for the basic BGP example involves the company network being connected to both ISPs as shown below. In this configuration the FortiGate unit is the BGP border router between the Company AS, ISP1's AS, and ISP2's AS.

The components of the layout include:

- The Company AS (AS number 1) is connected to ISP1 and ISP2 through the FortiGate unit.
- The Company has one internal network — the Head Office network at 10.11.101.0/24.
- The FortiGate unit internal interface is on the the Company internal network with an IP address of 10.11.101.110.
- The FortiGate unit external1 interface is connected to ISP1's network with an IP address of 172.20.111.5, an address supplied by the ISP.
- The FortiGate unit external2 interface is connected to IPS2's network with an IP address of 172.20.222.5, an address supplied by the ISP.
- ISP1 AS has an AS number of 650001, and ISP2 has an AS number of 650002.
- Both ISPs are connected to the Internet.
- The ISP1 border router is a neighbor (peer) of the FortiGate unit. It has an address of 172.21.111.4.
- The ISP2 border router is a neighbor (peer) of the FortiGate unit. It has an address of 172.22.222.4.
- Apart from graceful restart, and shorter timers (holdtimer, and keepalive) default settings are to be used whenever possible.

Figure 11:Basic BGP network topology



Assumptions

The basic BGP configuration procedure follows these assumptions:

- ISP1 is the preferred route, and ISP2 is the secondary route
- all basic configuration can be completed in both GUI and CLI
- only one AS is used for the Company

For these reasons this example configuration does not include:

- [Bi-directional forwarding detection \(BFD\)](#)
- [Route maps](#)
- [Access lists](#)
- changing redistribution defaults — make link when example is set up
- IPv6

For more information on these features, see the corresponding section.

Configuring the FortiGate unit

In this topology, the FortiGate unit is the link between the Company Network and the ISP network. The FortiGate unit is the only BGP router on the Company Network, but there is at least one other BGP router on the ISP Network — there may be more but we don't have that information.

As mentioned in the general configuration steps, the ISP must be notified of the Company's BGP router configuration when complete as it will need to add the FortiGate BGP router as a neighbor router on its domain. This step is required for the FortiGate unit to receive BGP routing updates from the ISP network and outside networks.

If the ISP has any special BGP features enabled such as graceful restart, or route dampening that should be determined up front so those features can be enabled on the FortiGate unit.

To configure the FortiGate unit as a BGP router

1. [Configure interfaces and default routes](#)
2. [Configure firewall services, addresses, and policies](#)
3. [Set the FortiGate BGP information](#)
4. [Add the internal network to the AS](#)
5. [Additional FortiGate BGP configuration](#)

Configure interfaces and default routes

The FortiGate unit is connected to three networks — Company Network on the internal interface, ISP1 Network on external1 interface, and ISP2 on external2 interface.

This example uses basic interface settings. Check with your ISP to determine if additional settings are required such as setting the maximum MTU size, or if gateway detection is supported.

High end FortiGate units do not have interfaces labeled Internal, or External. Instead, for clarity's sake, we are using the alias feature to name interfaces for these roles.

Default routes to both external interfaces are configured here as well. Both are needed in case one goes offline. ISP1 is the primary connection and has a smaller administrative distance so it will be preferred over ISP2. Both distances are set low so they will be preferred over any learned routes.

To configure the FortiGate interfaces - web-based manager

1. Go to *System > Network > Interface*.
2. Edit port 1 (internal) interface.

3. Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.11.101.110/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Company internal network
Administrative Status	Up

4. Edit port 2 (external1) interface.
5. Set the following information, and select *OK*.

Alias	external1
IP/Netmask	172.21.111.5/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	ISP1 External BGP network
Administrative Status	Up

6. Edit port 3 (external2) interface.
7. Set the following information, and select *OK*.

Alias	external2
IP/Netmask	172.22.222.5/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	ISP2 External BGP network
Administrative Status	Up

To configure the FortiGate interfaces - CLI

```
config system interface
  edit port1
    set alias internal
    set ip 10.11.101.110 255.255.255.0
    set allowaccess http https ssh
    set description "Company internal network"
    set status up
  next
  edit port2
    set alias external1
    set ip 172.21.111.5 255.255.255.0
    set allowaccess https ssh
    set description "ISP1 External BGP network"
    set status up
  next
  edit port3
    set alias external2
    set ip 172.22.222.5 255.255.255.0
    set allowaccess https ssh
    set description "ISP2 External BGP network"
    set status up
  next
end
```

To configure default routes for both ISPs - web-based manager

1. Go to *Router > Static > Static Routes*.
2. Delete any existing routes with a IP/Mask of address of 0.0.0.0/0.0.0.0
3. Select *Create New*, and set the following information.

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port2
Gateway	172.21.111.5
Distance	10

4. Select OK.
5. Select *Create New*, and set the following information.

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port3
Gateway	172.22.222.5
Distance	15

6. Select OK.

To configure default routes for both ISPs - CLI

```
config router static
  edit 1
    set device "port2"
    set distance 10
    set gateway 172.21.111.5
  next
  edit 2
    set device "port3"
    set distance 15
    set gateway 172.22.222.5
  next
end
```

Configure firewall services, addresses, and policies

To create the security policies, first you must create the firewall services group that will include all the services that will be allowed, then you must define the addresses that will be used in the security policies, and lastly you configure the security policies themselves.

To keep the configuration simple, the Company is allowing only HTTP traffic out of the local network. This will allow employees access to the Internet and their web-mail. DNS services will also be allowed through the firewall.

The security policies will allow HTTP traffic (port 80 and port 8080), HTTPS traffic (port 443), FTP traffic (port 21), and DNS traffic (port 53 and port 953) in both directions. Also BGP (port 179) may need access through the firewall.



For added security, you may want to define a smaller range of addresses for the internal network. For example if only 20 addresses are used, only allow those addresses in the range.

In the interest of keeping things simple, a zone will be used to group the two ISP interfaces together. This will allow using one security policy to apply to both ISPs at the same time. Remember to block intra-zone traffic as this will help prevent one ISP sending traffic to the other ISP through your FortiGate unit using your bandwidth. The zone keeps configuration simple, and in the future if there is a need for separate policies for each ISP, they can be created and the zone can be deleted.

The addresses that will be used are the addresses of the FortiGate unit internal and external ports, and the internal network.

More policies or services can be added in the future as applications are added to the network. For more information on security policies, see the firewall chapter of the [FortiGate Administration Guide](#).



When configuring security policies always enable logging to help you track and debug your traffic flow.

To create a firewall services group - web-based manager

1. Go to *Firewall Objects > Service > Groups*, and select *Create New*.

2. For *Group Name*, enter “Basic_Services”.
3. From *Available Services*, move the following six services over to the *Member* list — BGP, FTP, FTP_GET, FTP_PUT, DNS, HTTP, and HTTPS.
4. Select OK.

To create a firewall services group - CLI

```
config firewall service group
  edit "Basic_Services"
    set member "BGP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP" "HTTPS"
  next
end
```

To create a zone for the ISP interfaces - web-based manager

1. Go to *System > Network > Interfaces*.
2. Select the caret to the right of *Create New* and then select *Zone*.
3. Enter the following information.

Zone Name	ISPs
Block Intra-zone traffic	enable
interface members	port2 port3

4. Select OK.

To create a zone for the ISP interfaces - CLI

```
config system zone
  edit "ISPs"
    set interface "port2" "port3"
    set intrazone block
  next
end
```

To add the firewall addresses - web-based manager

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, and set the following information.

Address Name	Internal_network
Type	Subnet / IP Range
Subnet / IP Range	10.11.101.0 255.255.255.0
Interface	port1

3. Select OK.

To add the firewall addresses - CLI

```
config firewall address
  edit "Internal_network"
    set associated-interface "port1"
    set subnet 10.11.101.0 255.255.255.0
  next
end
```

To add the HTTP and DNS security policies - web-based manager

1. Go to *Policy > Policy > Policy*, and select *Create New*.
2. Set the following information.

Source Interface/Zone	port1(internal)
Source Address	Internal_network
Destination Interface/Zone	ISPs
Destination Address	All
Schedule	always
Service	Basic_services
Action	ACCEPT
Log Allowed Traffic	enable
Enable NAT	Enable
Comments	ISP1 basic services out policy

3. Select *OK*.
4. Select *Create New*, and set the following information.

Source Interface/Zone	ISPs
Source Address	all
Destination Interface/Zone	port1(internal)
Destination Address	Internal_network
Schedule	always
Service	Basic_services
Action	ACCEPT
Log Allowed Traffic	enable
NAT	Enable
Comments	ISP1 basic services in policy

To add the security policies - CLI

```
config firewall policy
  edit 1
    set srcintf "port1"
    set srcaddr "Internal_network"
    set dstintf "ISPs"
    set dstaddr "all"
    set schedule "always"
    set service "Basic_services"
    set action accept
    set nat enable
    set profile-status enable
    set logtraffic enable
    set comments "ISP1 basic services out policy"
  next
  edit 2
    set srcintf "ISPs"
    set srcaddr "all"
    set dstintf "port1"
    set dstaddr "Internal_network"
    set schedule "always"
    set service "Basic_services"
    set action accept
    set nat enable
    set profile-status enable
    set logtraffic enable
    set comments "ISP1 basic services in policy"
  next
end
```

Set the FortiGate BGP information

When using the default information, there are only two fields to set to configure the FortiGate unit as a BGP router.

For this configuration the FortiGate unit will be in a stub area with one route out — the ISP BGP router. Until you configure the ISP router as a neighbour, even that route out is not available. So while after this part of the configuration is complete your FortiGate unit will be running BGP, it won't know about any other routers running BGP until the next part of the configuration is complete.

To set the BGP router information - web-based manager

1. Go to *Router > Dynamic > BGP*.
2. Set the following information, and select OK.

Local AS	1
Router ID	10.11.101.110

To set the BGP router information - CLI

```
config router BGP
  set as 1
  set router-id 10.11.101.110
end
```

Add the internal network to the AS

The Company is one AS with the FortiGate unit configured as the BGP border router connecting that AS to the two ISPs ASes. The internal network in the Company's AS must be defined. If there were other networks in the company such as regional offices, they would be added here as well.

To set the networks in the AS - web-based manager

1. Go to *Router > Dynamic > BGP*.
2. In *Networks*, set the following information and select *OK*.

IP/Netmask	10.11.101.0/255.255.255.0
-------------------	---------------------------

To set the networks in the AS - CLI

```
config router bgp
  config network
  edit 1
    set prefix 10.11.101.0 255.255.255.0
  next
end
end
```

Add BGP neighbour information

The configuration will not work unless you set *Remote AS* neighbors. This can be done in either the web-based manager or the CLI.

To configure the BGP neighbors - web-based manager

1. Go to *Router > Dynamic > BGP*.
2. Add a *Neighbors IP* of 172.21.111.4 with the *Remote AS* set to 650001, then click *Add/Edit*.
3. Add another *Neighbors IP* of 172.22.222.4 with the *Remote AS* set to 650002, then click *Add/Edit*.

To configure the BGP neighbors - CLI

```
config router BGP
  set as 1
  config neighbor
  edit "172.21.111.4"
    set remote-as 650001
  next
  edit "172.22.222.4"
    set remote-as 650002
  next
end
end
```

Additional FortiGate BGP configuration

At this point that is all the settings that can be done in both the web-based manger and the CLI. The remaining configuration must be completed in the CLI.

These additional settings are mainly determined by your ISP requirements. They will determine your timers such as keep alive timers, if extended features like BFD and graceful restart are being used, and so on. For this example, some common simply features are being used to promote faster detections of network failures which will result in better service for the Company's internal network users.

The ISPs do not require authentication between peer routers.

These commands will enable or modify the following features on the FortiGate unit, and where possible on neighboring routers as well:

- `bestpath-med-missing-as-worst` — treats a route without an MED as the worst possible available route due to expected unreliability
- `fast-external-failover` — immediately reset the session information associated with BGP external peers if the link used to reach them goes down
- `graceful-restart*` — advertise reboots to neighbors so they do not see the router as offline, wait before declaring them offline, and how long to wait when they reboot before advertising updates. These commands applies to neighbors and are part of the BGP capabilities. This prevents unneeded routing updates.
- `holdtime-timer` — how long the router will wait for a keepalive message before declaring a router offline. A shorter time will find an offline router faster.
- `keepalive-timer` — how often the router sends out keepalive messages to neighbor routers to maintain those sessions.
- `log-neighbor-changes` — log changes to neighbor routers' status. This can be useful for troubleshooting from both internal and external networks.
- `connect-timer` — how long in seconds the FortiGate unit will try to reach this neighbor before declaring it offline.
- `weight` — used to prefer routes from one neighbor over the other. In this example ISP1 is the primary connection so it is weighted higher than ISP2

To configure additional BGP options - CLI

```
config router bgp
  set bestpath-med-missing-as-worst enable
  set fast-external-failover enable
  set graceful-restart enable
  set graceful-restart-time 120
  set graceful-stalepath-time 180
  set graceful-update-delay 180
  set holdtime-timer 120
  set keepalive-timer 45
  set log-neighbor-changes enable
config neighbor
  edit 172.21.111.4
    set connect-timer 60
    set description "ISP1"
    set holdtime-timer 120
    set keepalive-timer 45
    set weight 250
  next
  edit 172.22.222.4
    set connect-timer 60
    set description "ISP2"
    set holdtime-timer 120
    set keepalive-timer 45
    set weight 100
  next
end
end
```

Configuring other networking devices

There are two other networking devices that need to be configured: both ISPs' BGP routers.

The ISPs' routers must add the FortiGate unit as a neighbor so route updates can be sent in both directions. Note that ISP1 is not directly connected to ISP2 that we are aware of.

Inform both of your ISPs of your FortiGate unit's BGP information. Once they have configured their router, you can test your BGP connection to the Internet.

They will require your FortiGate unit's:

- IP address of the connected interface
- Router ID
- your Company's AS number

Testing this configuration

With the dual-homed BGP configuration in place, you should be able to send and receive traffic, send and receive routes, and not have any routing loops. Testing the networks will confirm things are working as expected.

In general for routing you need to look at the routing table on different routers to see what routes are being installed. You also need to sniff packets to see how traffic is being routed in real time. These two sources of information will normally tell you what you need to know.

Testing of this example's network configuration should be completed in two parts:

- Testing network connectivity
- Verifying the FortiGate unit's routing tables
- Verifying traffic routing
- Verifying the dual-homed side of the configuration

Testing network connectivity

A common first step in testing a new network topology is to test if you can reach the internet and other locations as you expect you should. If not, you may be prevented by cabling issues, software or other issues.

The easiest way to test connections is to use ping, once you ensure that all the FortiGate unit's interfaces and ISP routers have ping support enabled. Also ensure that the security policies allow ping through the firewall.

Connections to test in this example are the internal network to ISP1's router or the internet, and the same for ISP2. If you can connect on the external side of the Fortinet unit, try to ping the internal network. Those three tests should prove your basic network connections are working.



Once you have completed testing the network connectivity, turn off ping support on the external interfaces for additional security.

Verifying the FortiGate unit's routing tables

The FortiGate routing table contains the routes stored for future use. If you are expecting certain routes to be there and they are not, that is a good indicator that your configuration is not what you expected.

The CLI command `get router info routing-table details` will provide you with every route's routing protocol, destination address, gateway address, interface, weighting, and if the address is directly connected or not.

If you want to limit the display to BGP routes only, use the CLI command `get router info routing-table bgp`. If there are no BGP routes in the routing table, nothing will be displayed. In the CLI command you can replace BGP with static, or other routing protocols to only display those routes.

If you want to see the contents of the routing information database (RIB), use the CLI command `get router info routing-table database`. This will display the incoming routes that may or may not make it into the routing table.

Verifying traffic routing

Traffic may be reaching the internal network, but it may be using a different route than you think to get there.

Use a browser to try and access the Internet.

If needed, allow traceroute and other diag ports to be opened until things are working properly. Then remove access for them again.

Look for slow hops on the traceroute, or pings to a location, as they may indicate network loops that need to be fixed.

Any locations that have an unresolved traceroute or ping must be examined and fixed.

Use network packet sniffing to ensure traffic is being routed as you expect.

Verifying the dual-homed side of the configuration

Since there are two connections to the internet in this example, theoretically you can pull the plug on one of the ISP connections, and all traffic will go through the other connection. Alternately, you may choose to remove a default route to one ISP, remove that ISP's neighbor settings, or change the weightings to prefer other other ISP. These alternate ways to test dual-homing do not change physical cabling, which may be preferred in some situations.

If this does not work as expected, things to check include:

- default static routes — if these are wrong or don't exist, the traffic can't get out.
- BGP neighbor information — If the ISP router information is incorrect, the FortiGate unit won't be able to talk to it.

Redistributing and blocking routes in BGP

During normal BGP operation, peer routers redistribute routes from each other. However, in some specific situations it may be best to not advertise routes from one peer, such as if the peer is redundant with another peer (they share the same routes exactly), if it might be unreliable in some way, or some other reason. The FortiGate can also take routes it learns from other protocols and advertise them in BGP, for example OSPF or RIP. If your Company hosts its own web or email servers, external locations will require routes to your networks to reach those services.

In this example the Company has an internal network in an OSPF area, and is connected to a BGP AS and two BGP peers. Company goes through these two peers to reach the Internet. However, Peer 1 routes will not be advertised to Peer 2. The Company internal user and server networks are running OSPF, and will redistribute those routes to BGP so external locations can reach the web and email servers.

This section includes the following topics:

- [Network layout and assumptions](#)
- [Configuring the FortiGate unit](#)
- [Testing network configuration](#)

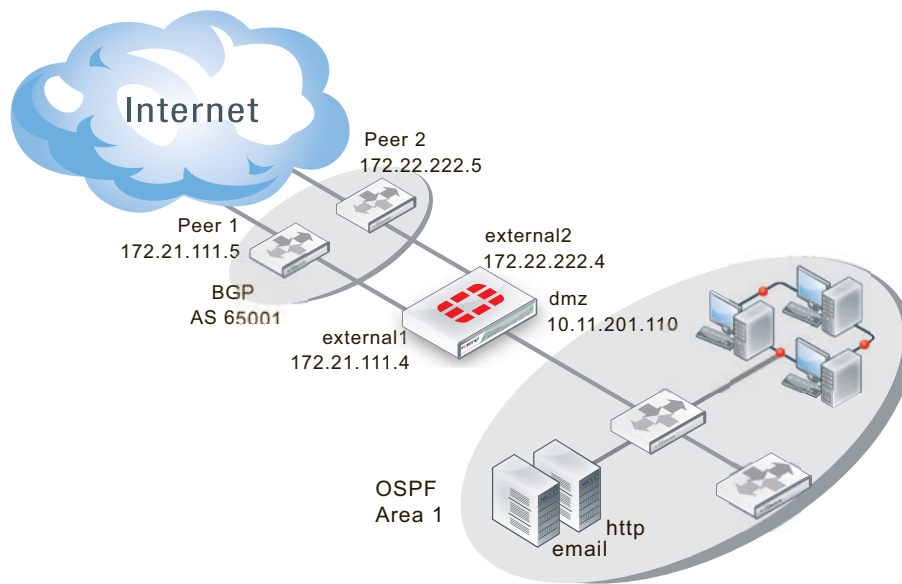
Network layout and assumptions

The network layout for the BGP redistributing routes example involves the company network being connected to two BGP peers as shown below. In this configuration the FortiGate unit is the BGP border router between the Company AS, and the peer routers.

The components of the layout include:

- There is only one BGP AS in this example — AS 65001, shared by the FortiGate unit and both peers.
- The Company's FortiGate unit connects to the Internet through two BGP peers.
- The Company internal networks on the dmz interface of the FortiGate unit with an IP of 10.11.201.0/24.
- The FortiGate units' interfaces are connected as follows:
 - port1 (dmz) has IP 10.11.201.110 and is the internal user and server network
 - port2 (external1) has IP 172.21.111.4 and is connected to Peer 1's network
 - port3 (external2) has IP 172.22.222.4 and is connected to Peer 2's network
- Peer 1 has IP 172.21.111.5, and Peer 2 has IP 172.22.222.5.
- OSPF Area 1 is configured on the dmz interface of the FortiGate unit, and is the routing protocol used by the internal users and servers.

Figure 12:BGP network topology



Assumptions

The the BGP redistributing routes configuration procedure follows these assumptions:

- the FortiGate unit has been configured following the Install Guide
- interfaces port1, port2, and port 3 exist on the FortiGate unit
- we don't know the router manufacturers of Peer 1 and Peer 2
- we don't know what other devices are on the BGP AS or OSPF Area
- all basic configuration can be completed in both GUI and CLI
- access lists and route maps will only be configured in CLI
- VDOMs are not enabled on the FortiGate unit

Configuring the FortiGate unit

1. Configuring the FortiGate unit — networks and firewalls
2. Configuring the FortiGate unit - BGP

3. [Configuring the FortiGate unit - OSPF](#)
4. [Configuring other networking devices](#)

Configuring the FortiGate unit – networks and firewalls

The FortiGate unit has three interfaces connected to networks – two external and one dmz.

Security policies must be in place to allow traffic to flow between these networks.

Firewall services will change depending on which routing protocol is being used on that network – either BGP or OSPF. Beyond that, all services that are allowed will be allowed in both directions due to the internal servers. The services allowed are web-server services (DNS, HTTP, HTTPS, SSH, NTP, FTP*, SYSLOG, and MYSQL), email services (POP3, IMAP, and SMTP), and general troubleshooting services (PING, TRACEROUTE). Those last two can be removed once the network is up and working properly to increase security. Other services can be added later as needed.

To configure the interfaces - GUI

1. Go to *System > Network > Interfaces*.
2. Edit port1 (dmz) interface.
3. Set the following information, and select *OK*.

Alias	dmz
IP/Netmask	10.11.201.110/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	OSPF internal networks
Administrative Status	Up

4. Edit port2 (external1) interface.
5. Set the following information, and select *OK*.

Alias	external1
IP/Netmask	172.21.111.4/255.255.255.0
Administrative Access	HTTPS SSH
Description	BGP external Peer 1
Administrative Status	Up

6. Edit port3 (external2) interface.
7. Set the following information, and select *OK*.

Alias	external2
IP/Netmask	172.22.222.4/255.255.255.0
Administrative Access	HTTPS SSH
Description	BGP external2 Peer2
Administrative Status	Up

To configure the FortiGate interfaces (CLI)

```
config system interface
  edit port1
    set alias dmz
    set ip 10.11.201.110 255.255.255.0
    set allowaccess https ssh ping
    set description "OSPF internal networks"
    set status up
  next
  edit port2
    set alias external1
    set ip 172.21.111.5 255.255.255.0
    set allowaccess https ssh
    set description "external1 Peer 1"
    set status up
  next
  edit port3
    set alias external2
    set ip 172.22.222.5 255.255.255.0
    set allowaccess https ssh
    set description "external2 Peer 2"
    set status up
  next
end
```

To configure the firewall addresses - GUI

1. Go to *Firewall Objects > Address > Addresses*.
2. Select *Create New*, and set the following information.

Address Name	BGP_services
Type	Subnet / IP Range
Subnet / IP Range	10.11.201.0 255.255.255.0
Interface	port1

3. Select *OK*.
4. Select *Create New*, and enter the following information:
5. Select *OK*.

To configure the firewall addresses - CLI

```
config firewall address
  edit "BGP_services"
    set associated-interface "port1"
    set subnet 10.11.201.0 255.255.255.0
  next
end
```

To configure firewall service groups - GUI

1. Go to *Firewall Objects > Service > Groups*.
2. Select *Create New*.
3. Name the group BGP_Services.
4. Move the following services to the right list: BGP, DNS, FTP, FTP_GET, FTP_PUT, HTTP, HTTPS, IMAP, MYSQL, NTP, PING, POP3, SMTP, SSH, SYSLOG, and TRACEROUTE.
5. Select *OK*.
6. Select *Create New*.
7. Name the group OSPF_Services.
8. Move the following services to the right list: DNS, FTP, FTP_GET, FTP_PUT, HTTP, HTTPS, IMAP, MYSQL, NTP, OSPF, PING, POP3, SMTP, SSH, SYSLOG, and TRACEROUTE.
9. Select *OK*.

To configure firewall service groups - CLI

```
config firewall service group
  edit "BGP_services"
    set member "BGP", "DHCP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP"
      "HTTPS" "IMAP" "MYSQL" "NTP" "PING" "POP3" "SMTP" "SSH"
      "TRACEROUTE" "SYSLOG"
  next
  edit "OSPF_services"
    set member "DHCP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP" "HTTPS"
      "IMAP" "MYSQL" "NTP" "PING" "POP3" "SMTP" "SSH"
      "TRACEROUTE" "SYSLOG" "OSPF"
  next
end
```

Configuring the FortiGate unit - BGP

The only change from the standard BGP configuration for this example is configuring the blocking Peer 1's routes from being advertised to Peer 2. From the network topology you can guess that both of these peers likely share many routes in common and it makes no sense to advertise unneeded routes.

Blocking Peer 1's routes to Peer 2 is done with distribute-list-out keyword. They allow you to select which routes you will advertise to a neighbor using an access list. In this case we will block all incoming routes from Peer 1 when we send updates to Peer 2. Otherwise Peer 1 and Peer 2 are regular neighbors.

The FortiGate unit will redistribute routes learned from OSPF into BGP.

This is advanced configuration and the commands are only available in the CLI.

To create access list to block Peer 1 - CLI

```
config access-list
  edit "block_peer1"
    config rule
      edit 1
        set prefix 172.21.111.0 255.255.255.0
        set action deny
        set exact-match enable
      end
    end
  end
end
```

To configure BGP on the FortiGate unit - CLI

```
config router bgp
  set as 65001
  set router-id 10.11.201.110
  config redistribute ospf
    set status enable
  end
  config neighbor
    edit 172.22.222.5
      set remote-as 65001
      set distribute-list-out "block_peer1"
    next
    edit 172.21.111.5
      set remote-as 65001
    end
  end
end
```

Configuring the FortiGate unit - OSPF

This configuration involves only one OSPF Area, so all traffic will be intra-area. If there were two or more areas with traffic going between them it would be inter-area traffic. These two types are comparable to BGP's traffic within one AS (iBGP) or between multiple ASes (eBGP). Redistributing routes from OSPF to BGP is considered external because either the start or end point is a different routing protocol.

The OSPF configuration is basic apart from redistributing BGP routes learned.

To configure OSPF on the FortiGate unit - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. For Router ID enter `10.11.201.110` and then select *Apply*.
3. Under *Advanced Options* and *Redistribute*, select *BGP* and set BGP metric to 1.
4. For *Areas*, select *Create New*, enter the following information and then select *OK*.

Area (IP)	0.0.0.0
Type	Regular
Authentication	None

5. For *Networks*, select *Create New*.

6. Enter 10.11.201.0/255.255.255.0 for *IP/Netmask*, and select *OK*.
7. For *Interfaces*, select *Create New*.
8. Enter `OSPF_dmz_network` for *Name*.
9. Select `port1(dmz)` for *Interface*, and then select *OK*.

To configure OSPF on the FortiGate unit - CLI

```
config router ospf
  set router-id 10.11.201.110
  config area
    edit 0.0.0.0
      set type regular
      set authentication none
    end
  config network
    edit 1
      set area 0.0.0.0
      set prefix 10.11.201.0 255.255.255.0
    end
  config interface
    edit "OSPF_dmz_network"
      set interface port1(dmz)
      set status enable
    end
  config redistribute bgp
    set status enable
    set metric 1
  end
end
```

Configuring other networking devices

As with all BGP configurations, the peer routers will need to be updated with the FortiGate unit's BGP information including IP address, AS number, and what capabilities are being used such as IPv6, graceful restart, BFD, and so on.

Testing network configuration

Testing this configuration involves the standard connectivity checks, but also ensuring that routes are being passed between protocols as expected.

Check the routing table on the FortiGate unit to ensure that routes from both OSPF and BGP are present.

Check the routing table on devices on the OSPF network for routes redistributed from BGP. Also check those devices for connectivity to the Internet.

Check the routing table on Peer 2 to ensure no routes from Peer 1 are present, but routes from the internal OSPF network are present.

For help with troubleshooting, see [“Troubleshooting BGP” on page 110](#).

Open Shortest Path First (OSPF)

This section describes OSPF routing.

The following topics are included in this section:

- OSPF Background and concepts
- Troubleshooting OSPF
- Basic OSPF example
- Advanced inter-area OSPF example
- Controlling redundant links by cost

OSPF Background and concepts

OSPF (Open Shortest Path First) is a link-state interior routing protocol, that is widely used in large enterprise organizations. It only routes packets within a single autonomous system (AS). This is different from BGP as BGP can communicate between ASes.

This section includes:

- Background
- The parts and terminology of OSPF
- How OSPF works

Background

OSPF version 2 was defined in 1998 in RFC 2328. OSPF was designed to support classless IP addressing, and variable subnet masks. This was a shortcoming of the earlier RIP protocols.

Updates to OSPF version 2 are included in OSPF version 3 defined in 2008 in RFC 5340. OSPF3 includes support for IPv6 addressing where previously OSPF2 only supports IPv4 addressing.

The main benefit of OSPF is that it detects link failures in the network quickly and within seconds has converged network traffic successfully without any networking loops. Also OSPF has many features to control which routes are propagated and which are not, maintaining smaller routing tables. OSPF can also provide better load-balancing on external links than other interior routing protocols.

The parts and terminology of OSPF

Parts and terminology of OSPF includes:

- OSPFv3 and IPv6
- Router ID
- Adjacency
- Designated router (DR) and backup router (BDR)
- Area
- Authentication
- Hello and dead intervals

OSPFv3 and IPv6

OSPFv3 (OSPF version 3) includes support for IPv6. Generally, all IP addresses are in IPv6 format instead of IPv4. However, OSPFv3 area numbers use the same 32-bit numbering system as OSPFv2, as described in [RFC 2740](#). Likewise, the router ID and area ID are in the same format as OSPFv2.

As with most advanced routing features on your FortiGate unit, IPv6 settings for dynamic routing protocols must be enabled before they will be visible in the GUI. To enable IPv6 configuration in the GUI, enable it in *System > Admin > Settings*.

For IPv6, the main difference in OSPFv3 is that, rather than using a network statement to enable OSPFv3 on an interface, you define OSPF6 (OSPF for IPv6) interfaces, which are bound to interface and area. This configuration must be done in the CLI, as follows (with sample interfaces and addresses):

```
config router ospf6
  config area
    edit 0.0.0.0
  next
end
config ospf6-interface
  edit "tunnel"
  set interface "to_FGT300A-7"
next
  edit "internal_lan"
  set interface "port1"
next
  set router-id 10.174.0.113
end
```

Note that OSPFv3 neighbors use link-local IPv6 addresses, but with broadcast and point-to-point network types, neighbors are automatically discovered. You only have to manually configure neighbours when using non-broadcast network types.

Router ID

In OSPF, each router has a unique 32-bit number called its Router ID. Often this 32-bit number is written the same as a 32-bit IPv4 address would be written in dotted decimal notation. However some brands of routers, such as Cisco routers, support a router ID entered as an integer instead of an IP address.

It is a good idea to not use IP address in use on the router for the router ID number. The router ID does not have to be a particular IP address on the router. By choosing a different number, it will be harder to get confused which number you are looking at. A good idea can be to use the as much of the area's number as possible. For example if you have 15 routers in area 0.0.0.0 they could be numbered from 0.0.0.1 to 0.0.0.15. If you have an area 1.1.1.1, then routers in that area could start at 1.1.1.10 for example.

You can manually set the router ID on your FortiGate unit.

To manually set an OSPF router ID of 0.0.1.1 - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. For *Router ID*, enter 0.0.1.1.
3. Select *Apply*.

To manually set an OSPF router ID of 0.0.1.1 - CLI

```
config router ospf
  set router-id 0.0.1.1
end
```

Adjacency

In an OSPF routing network, when an OSPF router boots up it sends out OSPF Hello packets to find any neighbors, routers that have access to the same network as the router booting up. Once neighbors are discovered and Hello packets are exchanged, updates are sent, and the Link State databases of both neighbors are synchronized. At this point these neighbors are said to be adjacent.

For two OSPF routers to become neighbors, the following conditions must be met.

- The subnet mask used on both routers must be the same subnet.
- The subnet number derived using the subnet mask and each router's interface IP address must match.
- The Hello interval & The Dead interval must match.
- The routers must have the same OSPF area ID. If they are in different areas, they are not neighbors.
- If authentication is used, they must pass authentication checks.

If any of these parameters are different between the two routers, the routers do not become OSPF neighbors and cannot be adjacent. If the routers become neighbors, they are adjacent.

Adjacency and neighbors

Neighbor routers can be in a Two-Way state, and not be adjacent. Adjacent routers normally have a neighbour state of FULL. Neighbors only exchange Hello packets, and do not exchange routing updates. Adjacent routers exchange LSAs (LSDB information) as well as Hello packets. A good example of an adjacent pair of routers is the DR and BDR.

You can check on the state of an OSPF neighbor using the CLI command `get router info ospf neighbor all`. See [“Checking the state of OSPF neighbors” on page 149](#).

Why adjacency is important

It is important to have adjacent pairs of routers in the OSPF routing domain because routing protocol packets are only passed between adjacent routers. This means adjacency is required for two OSPF routers to exchange routes. If there is no adjacency between two routers, such as one on the 172.20.120.0 network and another on the 10.11.101.0 network, the routers do not exchange routes. This makes sense because if all OSPF routers on the OSPF domain exchanged updates it would flood the network. Also, it is better for updates to progress through adjacent routers to ensure there are no outages along the way. Otherwise, updates could skip over routers that are potentially offline, causing longer routing outages and delays while the OSPF domain learns of this outage later on.

If the OSPF network has multiple border routers and multiple connections to external networks, the designated router (DR) determines which router pairs become adjacent. The DR can accomplish this because it maintains the complete topology of the OSPF domain, including which router pairs are adjacent. The BDR also has this information in case the DR goes offline.

Designated router (DR) and backup router (BDR)

In OSPF a router can have a number of different roles to play.

A designated router (DR) is the designated broadcasting router interface for an AS. It looks after all the initial contact and other routing administration traffic. Having only one router do all this greatly reduces the network traffic and collisions.

If something happens and the designated router goes offline, the backup designated router (BDR) takes over. An OSPF FortiGate unit interface can become either a DR or BDR. Both the DR and the BDR cover the same area, and are elected at the same time. The election process doesn't have many rules, but the exceptions can become complex.

Benefits

The OSPF concept of the designated router is a big step above RIP. With all RIP routers doing their own updates all the time, RIP suffers from frequent and sometimes unnecessary updates that can slow down your network. With OSPF, not only do routing changes only happen when a link-state changes instead of any tiny change to the routing table, but the designated router reduces this overhead traffic even more.

However, smaller network topologies may only have a couple routers besides the designated router. This may seem excessive, but it maintains the proper OSPF form and it will still reduce the administration traffic but to a lesser extent than on a large network. Also, your network topology will be ready whenever you choose to expand your network.

DR and BDR election

An election chooses the DR and BDR from all the available routers. The election is primarily based on the priority setting of the routers—the highest priority becomes the DR, and the second highest becomes BDR. To resolve any ties, the router with the highest router ID wins. For example 192.168.0.1 would win over 10.1.1.2.

The router priority can vary from 0 to 255, but at 0 a router will never become a DR or BDR. If a router with a higher priority comes on line after the election, it must wait until after the DR and BDR go offline before it would become the DR.

If the original DR goes offline, but then is available when the BDR goes offline later on, the original DR will be promoted back to DR without an election leaving the new BDR as it is.

With your FortiGate unit, to configure the port1 interface to be a potential OSPF designated router or backup designated router called `ospf_DR` on the network, you need to raise the priority of the router to a very high number such as 250 out of 255. This will ensure the interface has a chance to be a DR, but will not guarantee that it will be one. Give the interface a low numbered IP address—such as 10.1.1.1 instead of 192.168.1.1—to help ensure it becomes a DR, but that is not part of this example. Enter the following command:

```
config router ospf
  config ospf-interface
    edit "ospf_DR"
      set priority 250
    end
  end
```

Area

An OSPF area is a smaller part of the larger OSPF AS. Areas are used to limit the link-state updates that are sent out. The flooding used for these updates would overwhelm a large network, so it is divided into these smaller areas for manageability.

Within an area if there are two or more routers that are viable, there will always be a designated router (DR) and a backup DR (BDR). For more on these router roles, see [“Designated router \(DR\) and backup router \(BDR\)” on page 138](#).

Defining a private OSPF area, involves:

- assigning a 32-bit number to the area that is unique on your network
- defining the characteristics of one or more OSPF areas
- creating associations between the OSPF areas that you defined and the local networks to include in the OSPF area
- if required, adjusting the settings of OSPF-enabled interfaces.



IPv6 OSPF area numbers use the same 32-bit number notation as IPv4 OSPF.

If you are using the web-based manager to perform these tasks, follow the procedures summarized below.

FortiGate units support the four main types of OSPF area:

- [Backbone area](#)
- [NSSA](#)
- [Stub area](#)
- [Regular area](#)

Backbone area

Every OSPF network has at least one AS, and every OSPF network has a backbone area. The backbone is the main area, or possibly the only area. All other OSPF areas are connected to a backbone area. This means if two areas want to pass routing information back and forth, that routing information will go through the backbone on its way between those areas. For this reason the backbone not only has to connect to all other areas in the network, but also be uninterrupted to be able to pass traffic to all points of the network.

The backbone area is referred to as area 0 because it has an IP address of 0.0.0.0.

Stub area

A stub area is an OSPF area that receives no outside routes advertised into it, and all routing in it is based on a default route. This essentially isolates it from outside areas.

Stub areas are useful for small networks that are part of a larger organization, especially if the networking equipment can't handle routing large amounts of traffic passing through, or there are other reasons to prevent outside traffic, such as security. For example most organizations don't want their accounting department to be the center of their network with everyone's traffic passing through there. It would increase the security risks, slow down their network, and it generally doesn't make sense.

A variation on the stub area is the totally stubby area. It is a stub area that does not allow summarized routes.

NSSA

A not-so-stubby-area (NSSA) is a stub area that allows for external routes to be injected into it. While it still does not allow routes from external areas, it is not limited to only using the default route for internal routing.

Regular area

A regular area is what all the other ASes are, all the non-backbone, non-stub, non-NSSA areas. A regular area generally has a connection to the backbone, does receive advertisements of outside routes, and does not have an area number of 0.0.0.0.

Authentication

In the OSPF packet header are two authentication related fields — AuType, and Authentication.

All OSPF packet traffic is authenticated. Multiple types of authentication are supported in OSPFv2. However in OSPFv3, there is no authentication built-in but it is assumed that IPsec will be used for authentication instead.

Packets that fail authentication are discarded.

Null authentication

Null authentication indicates there is no authentication being used. In this case the 16-byte Authentication field is not checked, and can be any value. However checksumming is still used to locate errors. On your FortiGate this is the `none` option for authentication.

Simple Password authentication

Simple password refers to a standard plain text string of characters. The same password is used for all transactions on a network. The main use of this type of authentication is to prevent routers from accidentally joining the network. Simple password authentication is vulnerable to many forms of attack, and is not recommended as a secure form of authentication.

Cryptographic authentication

Cryptographic authentication involves the use of a shared secret key to authenticate all router traffic on a network. The key is never sent over the network in the clear—a packet is sent and a condensed and encrypted form of the packet is appended to the end of the packet. A non-repeating sequence number is included in the OSPF packet to protect against replay attacks that could try to use already sent packets to disrupt the network. When a packet is accepted as authentic the authentication sequence number is set to the packet sequence number. If a replay attack is attempted, the packet sent will be out of sequence and ignored.

Your FortiGate unit supports all three levels of authentication through the authentication keyword associated with creating an OSPF interface .

For example to create an OSPF interface called `Accounting` on the `port1` interface that is a broadcast interface, has a hello interval of 10 seconds, has a dead interval of 40 seconds, uses text authentication (simple password) with a password of “`ospf_test`”, enter the command:

```
config router ospf
  config ospf-interface
    edit Accounting
      set interface port1
      set network-type broadcast
      set hello-interval 10
      set dead-interval 40
      set authentication text
      set authentication-key "ospf_test"
    end
  end
```

Hello and dead intervals

The OSPF Hello protocol is used to discover and maintain communications with neighboring routers.

Hello packets are sent out at a regular interval for this purpose. The DR sends out the Hello packets. In a broadcast network, the multicast address of 224.0.0.5 is used to send out Hello packets. New routers on the network listen for and reply to these packets to join the OSPF area. If a new router never receives a Hello packet, other routers will not know it is there and will not communicate with it. However, once a new router is discovered the DR adds it to the list of routers in that area and it is integrated into the routing calculations.

Dead interval is the time other routers will wait before declaring a neighbor dead (offline). Setting a reasonable dead interval is very important. If this interval is too short, routers will be declared offline when they are just slow or momentarily inaccessible, and link-state updates will happen more than they need to, using more bandwidth. If the dead interval is too long, it will slow down network traffic overall if online routers attempt to contact offline ones instead of re-routing traffic. The CLI syntax for OSPF dead interval follows:

```
config ospf-interface
  edit ospf1
    set interface port1
    set network-type broadcast
    set dead-interval 1
end
```

Access Lists

Access lists are filters used by FortiGate unit OSPF routing. An access list provides a list of IP addresses and the action to take for them — essentially an access list makes it easy to group addresses that will be treated the same into the same group, independent of their subnets or other matching qualities. You add a rule for each address or subnet that you want to include, specifying the action to take for it. For example if you wanted all traffic from one department to be routed a particular way, even in different buildings, you can add all the addresses to an access list and then handle that list all at once.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or to match the prefix and any more specific prefix.

The FortiGate unit attempts to match a packet against the rules in an access list starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found the default action is deny.

Access lists greatly speed up configuration and network management. When there is a problem, you can check each list instead of individual addresses. Also, it eases troubleshooting since if all addresses on one list have problems, it eliminates many possible causes right away.

If you are using the OSPF+ IPv6 protocols you will need to use access-list6, the IPv6 version of access list. The only difference is that access-list6 uses IPv6 addresses.

For example, if you want to create an access list called `test_list` that only allows an exact match of `10.10.10.10` and `11.11.11.11`, enter the command:

```
config router access-list
  edit test_list
  config rule
    edit 1
      set prefix 10.10.10.10 255.255.255.255
      set action allow
      set exact-match enable
    next
    edit 2
      set prefix 11.11.11.11 255.255.255.255
      set action allow
      set exact-match enable
    end
  end
end
```

Another example is if you want to deny ranges of addresses in IPv6 that start with the IPv6 equivalents of `10.10.10.10` and `11.11.11.11`, enter the command `access-list6` as follows:

```
config router access-list6
  edit test_list_ip6
  config rule
    edit 1
      set prefix6 2002:A0A:A0A:0:0:0:0:0/48
      set action deny
    next
    edit 2
      set prefix6 2002:B0B:B0B:0:0:0:0:0/48
      set action deny
    end
  end
```

To use an `access_list`, you must call it from a routing protocol such as RIP. The following example uses the `access_list` from the earlier example called `test_list` to match routes coming in on the `port1` interface. When there is a match, it will add 3 to the hop count metric for those routes to artificially decrease their priority. Enter the following command:

```
config router ospf
  config distribute-list
    edit 5
      set access-list test_list
      set protocol connected
    end
```

If you are setting a prefix of `128.0.0.0`, use the format `128.0.0.0/1`. The default route `0.0.0.0/0` can not be exactly matched with an access-list. A prefix-list must be used for this purpose.

How OSPF works

An OSPF installation consists of one or more areas. An OSPF area is typically divided into logical areas linked by Area Border Routers. A group of contiguous networks form an area. An Area Border Router (ABR) links one or more areas to the OSPF network backbone (area ID 0). See [“Area border router \(ABR\)” on page 55](#).

OSPF is an interior routing protocol. It includes a backbone AS, and possibly additional ASes. The DR and BDR are elected from potential routers with the highest priorities. The DR handles much of the administration to lower the network traffic required. New routers are discovered through hello packets sent from the DR using the multicast address of 224.0.0.5. If the DR goes offline at any time, the BDR has a complete table of routes that it uses when it takes over as the DR router.

OSPF does not use UDP or TCP, but is encapsulated directly in IP datagrams as protocol 89. This is in contrast to RIP, or BGP. OSPF handles its own error detection and correction functions.

The OSPF protocol, when running on IPv4, can operate securely between routers, optionally using a variety of authentication methods to allow only trusted routers to participate in routing. OSPFv3, running on IPv6, no longer supports protocol-internal authentication. Instead, it relies on IPv6 protocol security (IPsec).

Other important parts of how OSPF works includes:

- [OSPF router discovery](#)
- [How OSPF works on FortiGate units](#)
- [External routes](#)
- [Link-state Database \(LSDB\) and route updates](#)
- [OSPF packets](#)

OSPF router discovery

OSPF-enabled routers generate Link-State Advertisements (LSA) and send them to their neighbors whenever the status of a neighbor changes or a new neighbor comes online. As long as the OSPF network is stable, LSAs between OSPF neighbors do not occur. An LSA identifies the interfaces of all OSPF-enabled routers in an area, and provides information that enables OSPF-enabled routers to select the shortest path to a destination. All LSA exchanges between OSPF-enabled routers are authenticated.

When a network of OSPF routers comes online, the follow steps occur.

1. When OSPF routers come online, they send out Hello packets to find other OSPF routers on their network segment.
2. When they discover other routers on their network segment, generally they become adjacent. Adjacent routers can exchange routing updates. See [“Adjacency” on page 138](#).
3. A DR and BDR are elected from the available routers using priority settings, and router ID. See [“Designated router \(DR\) and backup router \(BDR\)” on page 138](#), and [“DR and BDR election issues” on page 150](#).
4. Link state updates are sent between adjacent routers to map the topology of the OSPF area.
5. Once complete, the DR floods the network with the updates to ensure all OSPF routers in the area have the same OSPF route database. After the initial update, there are very few required updates if the network is stable.

How OSPF works on FortiGate units

When a FortiGate unit interface is connected to an OSPF area, that unit can participate in OSPF communications. FortiGate units use the OSPF Hello protocol to acquire neighbors in an area. A neighbor is any router that is directly connected to the same area as the FortiGate unit, and ideally is adjacent with a state of Full. After initial contact, the FortiGate unit exchanges Hello packets with its OSPF neighbors regularly to confirm that the neighbors can be reached.

The number of routes that a FortiGate unit can learn through OSPF depends on the network topology. A single unit can support tens of thousands of routes if the OSPF network is configured properly.

External routes

OSPF is an internal routing protocol. OSPF external routes are routes where the destination is using a routing protocol other than OSPF. OSPF handles external routes by adjusting the cost of the route to include the cost of the other routing protocol. There are two methods of calculating this cost, used for OSPF E1 and OSPF E2.

OSPF external1 (E1)

In OSPF E1 the destination is outside of the OSPF domain. This requires a different metric to be used beyond the normal OSPF metrics. The new metric of a redistributed route is calculated by adding the external cost and the OSPF cost together.

OSPF external2 (E2)

OSPF E2 is the default external type when routes are redistributed outside of OSPF. With OSPF E2, the metric of the redistributed route is equivalent to the external cost only, expressed as an OSPF cost. Dropping the OSPF portion can be useful in a number of situations, on border routers that have no OSPF portion for example or where the OSPF routing cost is negligible compared to the external routing cost.

Comparing E1 and E2

The best way to understand OSPF E1 and E2 routes is to check routing tables on OSPF routers. If you look at the routes on an OSPF border router, the redistributed routes will have an associated cost that represents only the external route, as there is no OSPF cost to the route due to it already being on the edge of the OSPF domain. However, if you look at that same route on a different OSPF router inside the OSPF routing domain, it will have a higher associated cost - essentially the external cost plus the cost over the OSPF domain to that border router. The border router uses OSPF E2, where the internal OSPF router uses OSPF E1 for the same route.

Viewing external routes

When you are trying to determine the costs for routes in your network to predict how traffic will be routed, you need to see the external OSPF routes and their associated costs. On your FortiGate unit, you find this information through your CLI.

To view external routes - CLI

You can view the whole routing table using `get router info routing-table all` to see all the routes including the OSPF external routes, or for a shorter list you can use the command `get router info routing-table ospf`. The letter at the left will be either E1 or E2 for external OSPF routes. The output of will look similar to the following, depending on what routes are in your routing table.

```
FGT620B# get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default

O*E2  0.0.0.0/0 [110/10] via 10.1.1.3, tunnel_wan2, 00:02:11
O     10.0.0.1/32 [110/300] via 10.1.1.3, tunnel_wan2, 00:02:11
S     0.0.0.0/0 [10/0] via 192.168.183.254, port2
S     1.0.0.0/8 [10/0] via 192.168.183.254, port2
```

Link-state Database (LSDB) and route updates

OSPF is based on links. The links between adjacent neighbor routers allow updates to be passed along the network. Network links allow the DR to flood the area with Link-state database (LSDB) updates. External links allow the OSPF area to connect to destinations outside the OSPF autonomous system. Information about these links is passed throughout the OSPF network as link-state updates.

The LSDB contains the information that defines the complete OSPF area, but the LSDB is not the routing table. It contains the information from all the link-state updates passed along the network. When there are no more changes required, and the network is stable then the LSDB on each router in the network will be the same. The DR will flood the LSDB to the area to ensure each router has the same LSDB.

To calculate the best route (shortest path) to a destination, the FortiGate unit applies the Shortest Path First (SPF) algorithm, based on Dijkstra's algorithm, to the accumulated link-state information. OSPF uses relative path cost metric for choosing the best route. The path cost can be any metric, but is typically the bandwidth of the path, how fast traffic will get from one point to another.

The path cost, similar to "distance" for RIP, imposes a penalty on the outgoing direction of a FortiGate unit interface. The path cost of a route is calculated by adding together all of the costs associated with the outgoing interfaces along the path to the destination. The lowest overall path cost indicates the best route, and generally the fastest route. Some brands of OSPF routers, such as Cisco, implement cost as a direct result of bandwidth between the routers. Generally this is a good cost metric because larger bandwidth means more traffic can travel without slowing down. To achieve this type of cost metric on FortiGate units, you need to set the cost for each interface manually in the CLI.



The inter-area routes may not be calculated when a Cisco type ABR has no fully adjacent neighbor in the backbone area. In this situation, the router considers summary-LSAs from all Actively summary-LSAs from all Actively Attached areas (RFC 3509).

The FortiGate unit dynamically updates its routing table based on the results of the SPF calculation to ensure that an OSPF packet will be routed using the shortest path to its destination. Depending on the network topology, the entries in the FortiGate unit routing table may include:

- the addresses of networks in the local OSPF area (to which packets are sent directly)
- routes to OSPF area border routers (to which packets destined for another area are sent)
- if the network contains OSPF areas and non-OSPF domains, routes to area boundary routers, which reside on the OSPF network backbone and are configured to forward packets to destinations outside the OSPF AS.

OSPF Route updates

Once the OSPF domain is established, there should be few updates required on a stable network. When updates occur and a decision is required concerning a new route, this is the general procedure.

1. Our router gets a new route, and needs to decide if it should go in the routing table.
2. The router has an up to date LSDB of the entire area, containing information about each router, the next hop to it, and most importantly the cost to get there.
3. Our router, turns the LSDB into a shortest path first (SPF) tree using Dijkstra's algorithm. It doesn't matter if there is more than one path to a router on the network, the SPF tree only cares about the shortest path to that router.

4. Once the SPF tree has been created, and shows the shortest paths to all the OSPF routers on the network, the work is done. If the new route is the best route, it will be part of that tree. If it is not the shortest route, it will not be included in the LSDB.
5. If there has been a change from the initial LSDB to the new SPF tree, a link state update will be sent out to let the other routers know about the change so they can update their LSDBs as well. This is vital since all routers on the OSPF area must have the same LSDB.
6. If there was no change between the LSDB and the SPF tree, no action is taken.

OSPF packets

Every OSPF packet starts with a standard 24-byte header, and another 24 bytes of information or more. The header contains all the information necessary to determine whether the packet should be accepted for further processing.

Table 11:OSPF packet

1-byte Version field	1-byte Type field	2-byte Packet length	3-byte Router ID
4-byte Area ID	2-byte Checksum	2-byte Auth Type	8-byte Authentication
4-byte Network Mask	2-byte Hello interval	1-byte Options field	1-byte Router Priority
4-byte Dead Router interval	4-byte DR field	4-byte BDR field	4-byte Neighbor ID

The following descriptions summarize the OSPF packet header fields.

Version field— The OSPF version number. This specification documents version 2 of the protocol.

Type field— There are 5 OSPF packet types. From one to five, respectively, they are Hello, Database Description, Link State Request, Link State Update, and Link State Acknowledgment.

Packet length— The length of the OSPF protocol packet in bytes. This length includes the standard OSPF 24-byte header, so all OSPF packets are at 24-bytes long.

Router ID— The Router ID of the packet's source.

Area ID— A 32-bit number identifying the area that this packet belongs to. All OSPF packets are associated with a single area. Most travel a single hop only. Packets travelling over a virtual link are labelled with the backbone Area ID of 0.0.0.0.

Checksum— The standard IP checksum of the entire contents of the packet, starting with the OSPF packet header but excluding the 64-bit authentication field. This checksum is calculated as the 16-bit one's complement of the one's complement sum of all the 16-bit words in the packet, excepting the authentication field. If the packet's length is not an integral number of 16-bit words, the packet is padded with a byte of zero before checksumming. The checksum is considered to be part of the packet authentication procedure; for some authentication types the checksum calculation is omitted.

Auth Type— Identifies the authentication procedure to be used for the packet. Authentication types include Null authentication (0), Simple password (1), Cryptographic authentication (2), and all others are reserved for future use.

Authentication— A 64-bit field for use by the authentication scheme. When AuType indicates no authentication is being used, the Authentication fields is not checked and can be any value.

When AuType is set to 2 (Cryptographic authentication), the 64-bit authentication field is split into the following four fields: Zero field, Key ID field, Authentication data length field, and Cryptographic sequence field.

The Key ID field indicates the key and algorithm used to create the message digest appended to the packet. The authentication data length field indicates how many bytes long the message digest is, and the cryptographic sequence number is at non-decreasing number that is set when the packet is received and authenticated to prevent replay attacks.

Network Mask—The subnet where this packet is valid.

Hello interval—The period of time between sending out Hello packets. See [“Hello and dead intervals” on page 142](#).

Options field— The OSPF protocol defines several optional capabilities. A router indicates the optional capabilities that it supports in its OSPF Hello packets, Database Description packets and in its LSAs. This enables routers supporting a mix of optional capabilities to coexist in a single Autonomous System.

Router priority—The priority between 0 and 255 that determines which routers become the DR and BDR. See [“Designated router \(DR\) and backup router \(BDR\)” on page 138](#).

Dead router interval—The period of time when there is no response from a router before it is declared dead. See [“Hello and dead intervals” on page 142](#).

DR and BDR fields—The DR and BDR fields each list the router that fills that role on this network, generally the routers with the highest priorities. See [“Designated router \(DR\) and backup router \(BDR\)” on page 138](#).

Neighbor ID—The ID number of a neighboring router. This ID is used to discover new routers and respond to them.

Troubleshooting OSPF

As with other dynamic routing protocols, OSPF has some issues that may need troubleshooting from time to time. For basic troubleshooting, see the FortiOS Handbook Troubleshooting chapter.

The more common issues include:

- [Clearing OSPF routes from the routing table](#)
- [Checking the state of OSPF neighbors](#)
- [Passive interface problems](#)
- [Timer problems](#)
- [Authentication issues](#)
- [DR and BDR election issues](#)

Clearing OSPF routes from the routing table

If you think the wrong route has been added to your routing table and you want to check it out, you first have to remove that route from your table before seeing if it is added back in or not. You can clear all or some OSPF neighbor connections (sessions) using the `execute router clear ospf` command. The `exec router clear` command is much more limiting for OSPF than it is for BGP. See [“Clearing routing table entries” on page 111](#).

For example, if you have routes in the OSPF routing table and you want to clear the specific route to IP address 10.10.10.1, you will have to clear all the OSPF entries. Enter the command:

```
execute router clear ospf process
```

Checking the state of OSPF neighbors

In OSPF each router sends out link state advertisements to find other routers on its network segment, and to create adjacencies with some of those routers. This is important because routing updates are only passed between adjacent routers. If two routers you believe to be adjacent are not, that can be the source of routing failures.

To identify this problem, you need to check the state of the OSPF neighbors of your FortiGate unit. Use the CLI command `get router info ospf neighbor all` to see all the neighbors for your FortiGate unit. You will see output in the form of:

```
FGT1 # get router info ospf neighbor
OSPF process 0:
Neighbor ID  Pri  State      Dead Time   Address      Interface
10.0.0.2     1   Full/ -   00:00:39   10.1.1.2    tunnel_wan1
10.0.0.2     1   Full/ -   00:00:34   10.1.1.4    tunnel_wan2
```

The important information here is the `State` column. Any neighbors that are not adjacent to your FortiGate unit will be reported in this column as something other than `Full`. If the state is `Down`, that router is offline.

Passive interface problems

A passive OSPF interface doesn't send out any updates. This means it can't be a DR, BDR, or an area border router among other things. It will depend on other neighbor routers to update its link-state table.

Passive interfaces can cause problems when they aren't receiving the routing updates you expect from their neighbors. This will result in the passive OSPF FortiGate unit interface having an incomplete or out-of-date link-state database, and it will not be able to properly route its traffic. It is possible that the passive interface is causing a hole in the network where no routers are passing updates to each other, however this is a rare situation.

If a passive interface is causing problems, there are simple methods to determine it is the cause. The easiest method is to make it an active interface, and if the issues disappear, then that was the cause. Another method is to examine the OSPF routing table and related information to see if it is incomplete compared to other neighbor routers. If this is the case, you can clear the routing table, reset the device and allow it to repopulate the table.

If you cannot make the interface active for some reason, you will have to change your network to fix the "hole" by adding more routers, or changing the relationship between the passive router's neighbors to provide better coverage.

Timer problems

A timer mismatch is when two routers have different values set for the same timer. For example if one router declares a router dead after 45 seconds and another waits for 4 minutes that difference in time will result in those two routers being out of synch for that period of time—one will still see that offline router as being online.

The easiest method to check the timers is to check the configuration on each router. Another method is to sniff some packets, and read the timer values in the packets themselves from different routers. Each packet contains the hello interval, and dead interval periods, so you can compare them easily enough.

Bi-directional Forwarding Detection (BFD)

Bi-directional Forwarding Detection (BFD) is a protocol used to quickly locate hardware failures in the network. Routers running BFD communicate with each other, and if a timer runs out on a

connection then that router is declared down. BFD then communicates this information to the routing protocol and the routing information is updated.

Authentication issues

OSPF has a number of authentication methods you can choose from. You may encounter problems with routers not authenticating as you expect. This will likely appear simply as one or more routers that have a blind spot in their routing - they won't acknowledge a router. This can be a problem if that router connects areas to the backbone as it will appear to be offline and unusable.

To confirm this is the issue, the easiest method is to turn off authentication on the neighboring routers. With no authentication between any routers, everything should flow normally.

Another method to confirm that authentication is the problem is to sniff packets, and look at their contents. The authentication type and password are right in the packets which makes it easy to confirm they are what you expect during real time. Its possible one or more routers is not configured as you expect and may be using the wrong authentication. This method is especially useful if there are a group of routers with these problems—it may only be one router causing the problem that is seen in multiple routers.

Once you have confirmed the problem is authentication related, you can decide how to handle it. You can turn off authentication and take your time to determine how to get your preferred authentication type back online. You can try another type of authentication, such as text instead of md5, which may have more success and still provide some level of protection. The important part is that once you confirm the problem, you can decide how to fix it properly.

DR and BDR election issues

You can force a particular router to become the DR and BDR by setting their priorities higher than any other OSPF routers in the area. This is a good idea when those routers have more resources to handle the traffic and extra work of the DR and BDR roles, since not all routers may be able to handle all that traffic.

However, if you set all the other routers to not have a chance at being elected, a priority of zero, you can run into problems if the DR and BDR go offline. The good part is that you will have some warning generally as the DR goes offline and the BDR is promoted to the DR position. But if the network segment with both the DR and BDR goes down, your network will have no way to send hello packets, send updates, or the other tasks the DR performs.

The solution to this is to always allow routers to have a chance at being promoted, even if you set their priority to one. In that case they would be the last choice, but if there are no other candidates you want that router to become the DR. Most networks would have already alerted you to the equipment problems, so this would be a temporary measure to keep the network traffic moving until you can find and fix the problem to get the real DR back online.

Basic OSPF example

This example sets up an OSPF network at a small office. There are 3 routers, all running OSPF v2. The border router connects to a BGP network.

All three routers in this example are FortiGate units. Router1 will be the designated router (DR) and router2 will be the backup DR (BDR) due to their priorities. Router3 will not be considered for either the DR or BDR elections. Instead, Router3 is the area border router (ASBR) routing all traffic to the ISP's BGP router on its way to the Internet.

Router2 has a modem connected that provides dialup access to the Internet as well, at a reduced bandwidth. This is a PPPoE connection to a DSL modem. This provides an alternate

route to the Internet if the other route goes down. The DSL connection is slow, and is charged by the amount of traffic. For these reasons OSPF will highly favor Router3's Internet access.

The DSL connection connects to an OSPF network with the ISP, so no redistribution of routes is required. The ISP network does have to be added to that router's configuration however.

This section includes the following topics:

- [Network layout and assumptions](#)
- [Configuring the FortiGate units](#)
- [Configuring OSPF on the FortiGate units](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

Network layout and assumptions

There are three FortiGate units acting as OSPF v2 routers on the network—Router1, Router2, and Router3. Router1 will be the designated router (DR), and Router 2 the BDR. Router3 is the area border router (ASBR) that connects to the external ISP router running BGP. Router2 has a PPPoE DSL connection that can access the Internet.

The Head Office network is connected to Router1 and Router2 on the 10.11.101.0 subnet.

Router1 and Router3 are connected over the 10.11.103.0 subnet.


Router2 and Router3 are connected over the 10.11.102.0 subnet.

The following table lists the router, interface, address, and role it is assigned.

Table 12:Routers, interfaces, and IP addresses for basic OSPF example network

Router name	Interface	IP address	Interface is connected to:
Router1 (DR)	Internal (port1)	10.11.101.1	Head office network, and Router2
	External (port2)	10.11.102.1	Router3
Router2 (BDR)	Internal (port1)	10.11.101.2	Head office network, and Router1
	External (port2)	10.11.103.2	Router3
	DSL (port3)	10.12.101.2	PPPoE DSL access
Router3 (ASBR)	Internal1 (port1)	10.11.102.3	Router1
	Internal2 (port2)	10.11.103.3	Router2
	External (port3)	172.20.120.3	ISP's BGP network

Figure 13:Basic OSPF network topology



Note that other subnets can be added to the internal interfaces without changing the configuration.

Assumptions

- The FortiGate units used in this example have interfaces named port1, port2, and port3.
- All FortiGate units in this example have factory default configuration with FortiOS 4.0 MR2 firmware installed, and are in NAT/Route operation mode.
- Basic firewalls are in place to allow unfiltered traffic between all connected interfaces in both directions.
- This OSPF network is not connected to any other OSPF networks.
- Both Internet connections are always available.
- The modem connection is very slow and expensive.
- Other devices may be on the network, but do not affect this basic configuration.
- Router3 is responsible for redistributing all routes into and out of the OSPF AS.

Configuring the FortiGate units

Each FortiGate unit needs the interfaces, and basic system information such as hostname configured.

This section includes:

- [Configuring Router1](#)
- [Configuring Router2](#)
- [Configuring Router3](#)

Configuring Router1

Router1 has two interfaces connected to the network—internal (port1) and external (port2). Its host name must be changed to Router1.

To configure Router1 interfaces - web-based manager

1. Go to *System > Dashboard > Status*.
2. Beside the host name, select *Change*.
3. Enter a hostname of `Router1`, and select *OK*.
4. Go to *System > Network > Interfaces*, edit port1, set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.11.101.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Head office and Router2
Administrative Status	Up

5. Edit port2, set the following information, and select *OK*.

Alias	External
IP/Netmask	10.11.102.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3
Administrative Status	Up

Configuring Router2

Router2 configuration is the same as Router1, except Router2 also has the DSL interface to configure.

The DSL interface is configured with a username of “user1” and a password of “ospf_example”. The default gateway will be retrieved from the ISP, and the defaults will be used for the rest of the PPPoE settings.

To configure Router2 interfaces - web-based manager

1. Go to *System > Dashboard > Status*.
2. Beside the host name, select *Change*.
3. Enter a hostname of `Router2`, and select *OK*.
4. Go to *System > Network > Interfaces*, edit port1, set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.11.101.2/255.255.255.0
Administrative Access	HTTPS SSH PING

Description	Head office and Router1
Administrative Status	Up

5. Edit port2, set the following information, and select *OK*.

Alias	External
IP/Netmask	10.11.103.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3
Administrative Status	Up

6. Edit DSL (port3), set the following information, and select *OK*.

Alias	DSL
Addressing Mode	PPPoE
Username	user1
Password	ospf_example
Unnumbered IP address	10.12.101.2/255.255.255.0
Retrieve default gateway from server	Enable
Administrative Access	HTTPS SSH PING
Description	DSL
Administrative Status	Up

Configuring Router3

Router3 is similar to Router1 and Router2 configurations. The main difference is the External (port3) interface connected to the ISP BGP network which has no administration access enabled for security reasons.

To configure Router3 interfaces - web-based manager

1. Go to *System > Status > Dashboard*.
2. Next to hostname, select *Change*.
3. Enter a hostname of `Router3`, and select *OK*.
4. Go to *System > Network > Interfaces*, edit port1, set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.11.102.3/255.255.255.0
Administrative Access	HTTPS SSH PING

Description	Router1
Administrative Status	Up

5. Edit port2, set the following information, and select *OK*.

Alias	Internal2
IP/Netmask	10.11.103.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router2
Administrative Status	Up

6. Edit port3, set the following information, and select *OK*.

Alias	External
IP/Netmask	172.20.120.3/255.255.255.0
Administrative Access	
Description	ISP BGP
Administrative Status	Up

Configuring OSPF on the FortiGate units

With the interfaces configured, now the FortiGate units can be configured for OSPF on those interfaces. All routers are part of the backbone 0.0.0.0 area, so there is no inter-area communications needed.

For a simple configuration there will be no authentication, no graceful restart or other advanced features, and timers will be left at their defaults. Also the costs for all interfaces will be left at 10, except for the modem and ISP interfaces where cost will be used to load balance traffic. Nearly all advanced features of OSPF are only available from the CLI.

The network that is defined covers all the subnets used in this example - 10.11.101.0, 10.11.102.0, and 10.11.103.0. All routes for these subnets will be advertised. If there are other interfaces on the FortiGate units that you do not want included in the OSPF routes, ensure those interfaces use a different subnet outside of the 10.11.0.0 network. If you want all interfaces to be advertised you can use an OSPF network of 0.0.0.0 .

Each router will configure:

- router ID
- area
- network
- two or three interfaces depending on the router
- priority for DR (Router1) and BDR (Router2)
- redistribute for ASBR (Router3)

This section includes:

- [Configuring OSPF on Router1](#)
- [Configuring OSPF on Router2](#)
- [Configuring OSPF on Router3](#)

Configuring OSPF on Router1

Router1 has a very high priority to ensure it becomes the DR for this area. Also Router1 has the lowest IP address to help ensure it will win in case there is a tie at some point. Otherwise it is a standard OSPF configuration. Setting the priority can only be done in the CLI, and it is for a specific OSPF interface.

To configure OSPF on Router1 - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. Set *Router ID* to `10.11.101.1` and select *Apply*.
3. In *Areas*, select *Create New*, set the following information, and select *OK*.

Area	0.0.0.0
Type	Regular
Authentication	none

4. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	10.11.0.0/255.255.0.0
Area	0.0.0.0

5. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Router1-Internal-DR
Interface	port1 (Internal)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

6. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Router1-External
Interface	port2 (External)
IP	0.0.0.0
Authentication	none

Timers (seconds)

Hello Interval	10
Dead Interval	40

7. Using the CLI, enter the following commands to set the priority for the Router1-Internal OSPF interface to maximum, ensuring this interface becomes the DR.

```
config router ospf
  config ospf-interface
    edit Router1-Internal-DR
      set priority 255
    end
```

To configure OSPF on Router1 - CLI

```
config router ospf
  set router-id 10.11.101.1
  config area
    edit 0.0.0.0
      next
    end
  config network
    edit 1
      set prefix 10.11.0.0/255.255.255.0
    next
  end
  config ospf-interface
    edit "Router1-Internal"
      set interface "port1"
      set priority 255
    next
    edit "Router1-External"
      set interface "port2"
    next
  end
end
```

Configuring OSPF on Router2

Router2 has a high priority to ensure it becomes the BDR for this area, and configures the DSL interface slightly differently—assume this will be a slower connection resulting in the need for longer timers, and a higher cost for this route.

Otherwise it is a standard OSPF configuration.

To configure OSPF on Router2 - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. Set *Router ID* to `10.11.101.2` and select *Apply*.

3. In *Areas*, select *Create New*, set the following information, and select *OK*.

Area	0.0.0.0
Type	Regular
Authentication	none

4. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	10.11.0.0/255.255.0.0
Area	0.0.0.0

5. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Router2-Internal
Interface	port1 (Internal)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

6. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Router2-External
Interface	port2 (External)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

7. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Router2-DSL
Interface	port3 (DSL)
IP	0.0.0.0
Authentication	none
Cost	50

Timers (seconds)

Hello Interval	20
Dead Interval	80

8. Using the CLI, enter the following commands to set the priority for the Router2-Internal OSPF interface to ensure this interface will become the BDR.

```
config router ospf
  config ospf-interface
    edit Router2-Internal
      set priority 250
    next
  end
```

To configure OSPF on Router2 - CLI

```
config router ospf
  set router-id 10.11.101.2
  config area
    edit 0.0.0.0
    next
  end
  config network
    edit 1
      set prefix 10.11.0.0/255.255.0.0
    next
  end
  config ospf-interface
    edit "Router2-Internal"
      set interface "port1"
      set priority 255
    next
    edit "Router2-External"
      set interface "port2"
    next
    edit "Router2-DSL"
      set interface "port3"
      set cost 50
    next
  end
end
```

Configuring OSPF on Router3

Router3 is more complex than the other two routers. The interfaces are straightforward, but this router has to import and export routes between OSPF and BGP. That requirement makes Router3 a border router or ASBR. Also Router3 needs a lower cost on its route to encourage all traffic to the Internet to route through it.

In the advanced OSPF options, Redistribute is enabled for Router3. It allows different types of routes, learned outside of OSPF, to be used in OSPF. Different metrics are assigned to these other types of routes to make them more or less preferred to regular OSPF routes.

To configure OSPF on Router3 - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. Set *Router ID* to `10.11.101.2` and select *Apply*.
3. Expand *Advanced Options*.
4. In *Redistribute*, set the following information, and select *OK*.

Route type	Redistribute	Metric
Connected	Enable	15
Static	Enable	15
RIP	Disable	n/a
BGP	Enable	5

5. In *Areas*, select *Create New*, set the following information, and select *OK*.

Area	0.0.0.0
Type	Regular
Authentication	none

6. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	10.11.0.0/255.255.0.0
Area	0.0.0.0

7. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Router3-Internal
Interface	port1 (Internal)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

8. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Router3-Internal2
Interface	port2 (Internal2)
IP	0.0.0.0

Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

9. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Router3-ISP-BGP
Interface	port3 (ISP-BGP)
IP	0.0.0.0
Authentication	none
Cost	2
Timers (seconds)	
Hello Interval	20
Dead Interval	80

10. Using the CLI, enter the following commands to set the priority for the Router3-Internal OSPF interface to ensure this interface will become the BDR.

```
config router ospf
  config ospf-interface
  edit Router3-Internal
    set priority 250
  next
end
```

To configure OSPF on Router3 - CLI

```
config router ospf
  set router-id 10.11.102.3
  config area
    edit 0.0.0.0
    next
  end
  config network
    edit 1
      set prefix 10.11.0.0/255.255.255.0
    next
    edit 2
      set prefix 172.20.120.0/255.255.255.0
    next
  end
  config ospf-interface
    edit "Router3-Internal"
      set interface "port1"
      set priority 255
    next
    edit "Router3-External"
      set interface "port2"
    next
    edit "Router3-ISP-BGP"
      set interface "port3"
      set cost 2
    next
  end
end
```

Configuring other networking devices

The other networking devices required in this configuration are on the two ISP networks, the BGP network for the main Internet connection, and the DSL backup connection.

In both cases, the ISPs need to be notified of the OSPF network settings including router IP addresses, timer settings, and so on. The ISP will use this information to configure its routers that connect to this OSPF network.

Testing network configuration

Testing the network configuration involves two parts: testing the network connectivity, and testing the OSPF routing.

To test the network connectivity use ping, traceroute, and other network tools.

To test the OSPF routing in this example, refer to the troubleshooting outlined in [“Troubleshooting OSPF” on page 148](#).

Advanced inter-area OSPF example

This example sets up an OSPF network at a large office. There are three areas, each with two routers. Typically OSPF areas would not be this small, and if they were the areas would be combined into one bigger area. However, the stub area services the accounting department which is very sensitive about their network and do not want any of their network information broadcast through the rest of the company. The backbone area contains the bulk of the company network devices. The regular area was established for various reasons such as hosting the company servers on a separate area with extra security.

One area is a small stub area that has no independent Internet connection, and only one connection to the backbone area. That connection between the stub area and the backbone area is only through a default route. No routes outside the stub area are advertised into that area. Another area is the backbone, which is connected to the other two areas. The third area has the Internet connection, and all traffic to and from the Internet must use that area's connection. If that traffic comes from the stub area, then that traffic is treating the backbone like a transit area that only uses it to get to another area.

In the stub area, a subnet of computers is running the RIP routing protocol and those routes must be redistributed into the OSPF areas.


This section includes the following topics:

- [Network layout and assumptions](#)
- [Configuring the FortiGate units](#)
- [Configuring OSPF on the FortiGate units](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

Network layout and assumptions

There are four FortiGate units in this network topology acting as OSPF routers.

Figure 14:Advanced inter-area OSPF network topology



Area 1.1.1.1 is a stub area with one FortiGate unit OSPF router called Router1 (DR). Its only access outside of that area is a default route to the backbone area, which is how it accesses the Internet. Traffic must go from the stub area, through the backbone, to the third area to reach the Internet. The backbone area in this configuration is called a transit area. Also in area 1.1.1.1 there is a RIP router that will be providing routes to the OSPF area through redistribution.

Area 0.0.0.0 is the backbone area, and has two FortiGate unit routers named Router2 (BDR) and Router3 (DR).

Area 2.2.2.2 is a regular area that has an Internet connection accessed by both the other two OSPF areas. There is only one FortiGate unit router in this area called Router4 (DR). This area is more secure and requires MD5 authentication by routers.

All areas have user networks connected, but they are not important for configuring the network layout for this example.

Internal interfaces are connected to internal user networks only. External1 interfaces are connected to the 10.11.110.0 network, joining Area 1.1.1.1 and Area 0.0.0.0.

External2 interfaces are connected to the 10.11.111.0 network, joining Area 0.0.0.0 and Area 2.2.2.2. The ISP interface is called ISP.

Table 13:Routers, areas, interfaces, IP addresses for advanced OSPF network

Router name	Area number and type	Interface	IP address
Router1 (DR)	1.1.1.1 - stub area (Accounting)	port1 (internal)	10.11.101.1
		port2 (external1)	10.11.110.1

Table 13:Routers, areas, interfaces, IP addresses for advanced OSPF network

Router2 (BDR)	0.0.0.0 - backbone area (R&D Network)	port1 (internal)	10.11.102.2
		port2 (external1)	10.11.110.2
		port3 (external2)	10.11.111.2
Router3 (DR)	0.0.0.0 - backbone area (R&D Network)	port1 (internal)	10.11.103.3
		port2 (external1)	10.11.110.3
		port3 (external2)	10.11.111.3
Router4 (DR)	2.2.2.2 - regular area (Network Admin)	port1 (internal)	10.11.104.4
		port2 (external2)	10.11.111.4
		port3 (ISP)	172.20.120.4

Note that other subnets can be added to the internal interfaces without changing the configuration.

Assumptions

- The FortiGate units used in this example have interfaces named port1, port2, and port3.
- All FortiGate units in this example have factory default configuration with FortiOS 4.0 MR2 firmware installed, and are in NAT/Route operation mode.
- During configuration, if settings are not directly referred to they will be left at default settings.
- Basic firewalls are in place to allow unfiltered traffic between all connected interfaces in both directions.
- This OSPF network is not connected to any other OSPF areas outside of this example.
- The Internet connection is always available.
- Other devices may be on the network, but do not affect this configuration.

Configuring the FortiGate units

This section configures the basic settings on the FortiGate units to be OSPF routers. These configurations include multiple interface settings, and hostname.

There are four FortiGate units in this example. The two units in the backbone area can be configured exactly the same except for IP addresses, so only router3 (the DR) configuration will be given with notes indicating router2 (the BDR) IP addresses.

Configuring the FortiGate units includes:

- [Configuring Router1](#)
- [Configuring Router2](#)
- [Configuring Router3](#)
- [Configuring Router4](#)

Configuring Router1

Router1 is part of the Accounting network stub area (1.1.1.1).

To configure Router1 interfaces - web-based manager

1. Go to *System > Dashboard > Status*.
2. Next to hostname, select *Change*.
3. Enter a hostname of `Router1`, and select *OK*.
4. Go to *System > Network > Interfaces*, edit port1, set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.11.101.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Accounting network
Administrative Status	Up

5. Edit port2, set the following information, and select *OK*.

Alias	External1
IP/Netmask	10.11.110.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Backbone network and internet
Administrative Status	Up

Configuring Router2

Router2 is part of the R&D network backbone area (0.0.0.0). Router2 and Router3 are in this area. They provide a redundant connection between area 1.1.1.1 and area 2.2.2.2.

Router2 has three interfaces configured; one to the internal network, and two to Router3 for redundancy.

To configure Router2 interfaces - web-based manager

1. Go to *System > Dashboard > Status*.
2. Next to hostname, select *Change*.
3. Enter a hostname of `Router2`, and select *OK*.
4. Go to *System > Network > Interfaces*, edit port1 (internal), set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.11.102.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

5. Edit port2 (external1), set the following information, and select *OK*.

Alias	external1
IP/Netmask	10.11.110.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3 first connection
Administrative Status	Up

6. Edit port3 (external2), set the following information, and select *OK*.

Alias	external2
IP/Netmask	10.11.111.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3 second connection
Administrative Status	Up

Configuring Router3

Router3 is part of the R&D network backbone area (0.0.0.0). Router2 and Router3 are in this area. They provide a redundant connection between area 1.1.1.1 and area 2.2.2.2.

To configure Router3 interfaces - web-based manager

1. Go to *System > Dashboard > Status*.
2. Next to hostname, select *Change*.
3. Enter a hostname of `Router3`, and select *OK*.
4. Go to *System > Network > Interfaces*, edit port1 (internal), set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.11.103.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

5. Edit port2 (external1), set the following information, and select *OK*.

Alias	external1
IP/Netmask	10.11.110.3/255.255.255.0
Administrative Access	HTTPS SSH PING

Description	Router2 first connection
Administrative Status	Up

6. Edit port3 (external2), set the following information, and select *OK*.

Alias	external2
IP/Netmask	10.11.111.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router2 second connection
Administrative Status	Up

Configuring Router4

Router4 is part of the Network Administration regular area (2.2.2.2). This area provides internet access for both area 1.1.1.1 and the backbone area.

This section configures interfaces and hostname.

To configure Router4 interfaces - web-based manager

1. Go to *System > Dashboard > Status*.
2. Next to hostname, select *Change*.
3. Enter a hostname of Router4, and select *OK*.
1. Go to *System > Network > Interfaces*.
2. Edit port1 (internal).
3. Set the following information, and select *OK*.

Alias	internal
IP/Netmask	10.11.101.4/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Accounting network
Administrative Status	Up

4. Edit port2 (external2).
5. Set the following information, and select *OK*.

Alias	external2
IP/Netmask	10.11.110.4/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Backbone and Accounting network
Administrative Status	Up

6. Edit port3 (ISP).

7. Set the following information, and select *OK*.

Alias	ISP
IP/Netmask	172.20.120.4/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	ISP and internet
Administrative Status	Up

Configuring OSPF on the FortiGate units

Three of the routers are designated routers (DR) and one is a backup DR (BDR). This is achieved through the lowest router ID numbers, or OSPF priority settings.

Also each area needs to be configured as each respective type of area - stub, backbone, or regular. This affects how routes are advertised into the area.

To configure OSPF on Router1 - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. Enter *10.11.101.1* for the *Router ID* and select *Apply*.
3. In *Areas*, select *Create New*, set the following information, and select *OK*.

Area	1.1.1.1
Type	Stub
Authentication	None

4. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	10.11.101.0/255.255.255.0
Area	1.1.1.1

5. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Accounting
Interface	port1 (internal)
IP	10.11.101.1
Authentication	None

6. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Backbone1
Interface	port2 (external1)
IP	10.11.110.1
Authentication	None

To configure OSPF on Router2 - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. Enter 10.11.102.2 for the *Router ID* and select *Apply*.
3. In *Areas*, select *Create New*, set the following information, and select *OK*.

Area	0.0.0.0
Type	Regular
Authentication	None

4. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	10.11.102.2/255.255.255.0
Area	0.0.0.0

5. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	10.11.110.2/255.255.255.0
Area	0.0.0.0

6. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	10.11.111.2/255.255.255.0
Area	0.0.0.0

7. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	RnD network
Interface	port1 (internal)
IP	10.11.102.2
Authentication	None

8. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Backbone1
Interface	port2 (external1)
IP	10.11.110.2
Authentication	None

9. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Backbone2
Interface	port3 (external2)

IP	10.11.111.2
Authentication	None

To configure OSPF on Router3 - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. Enter 10.11.103.3 for the *Router ID* and then select *Apply*.
3. In *Areas*, select *Create New*, set the following information, and then select *OK*.

Area	0.0.0.0
Type	Regular
Authentication	None

4. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	10.11.102.3/255.255.255.0
Area	0.0.0.0

5. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	10.11.110.3/255.255.255.0
Area	0.0.0.0

6. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	10.11.111.3/255.255.255.0
Area	0.0.0.0

7. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	RnD network
Interface	port1 (internal)
IP	10.11.103.3
Authentication	None

8. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Backbone1
Interface	port2 (external1)
IP	10.11.110.3
Authentication	None

9. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Backbone2
Interface	port3 (external2)
IP	10.11.111.3
Authentication	None

To configure OSPF on Router4 - web-based manager

1. Go to *Router > Dynamic > OSPF*.
2. Enter 10.11.104.4 for the *Router ID* and then select *Apply*.
3. In *Areas*, select *Create New*.
4. Set the following information, and select *OK*.

Area	2.2.2.2
Type	Regular
Authentication	None

5. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	10.11.104.0/255.255.255.0
Area	0.0.0.0

6. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	10.11.111.0/255.255.255.0
Area	0.0.0.0

7. In *Networks*, select *Create New*, set the following information, and select *OK*.

IP/Netmask	172.20.120.0/255.255.255.0
Area	0.0.0.0

8. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Network Admin network
Interface	port1 (internal)
IP	10.11.104.4
Authentication	None

9. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	Backbone2
Interface	port2 (external2)
IP	10.11.111.4
Authentication	None

10. In *Interfaces*, select *Create New*, set the following information, and select *OK*.

Name	ISP
Interface	port3 (ISP)
IP	172.20.120.4
Authentication	None

Configuring other networking devices

All network devices on this network are running OSPF routing. The user networks (Accounting, R&D, and Network Administration) are part of one of the three areas.

The ISP needs to be notified of your network configuration for area 2.2.2.2. Your ISP will not advertise your areas externally as they are intended as internal areas. External areas have assigned unique numbers. The area numbers used in this example are similar to the 10.0.0.0 and 192.168.0.0 subnets used in internal networking.

Testing network configuration

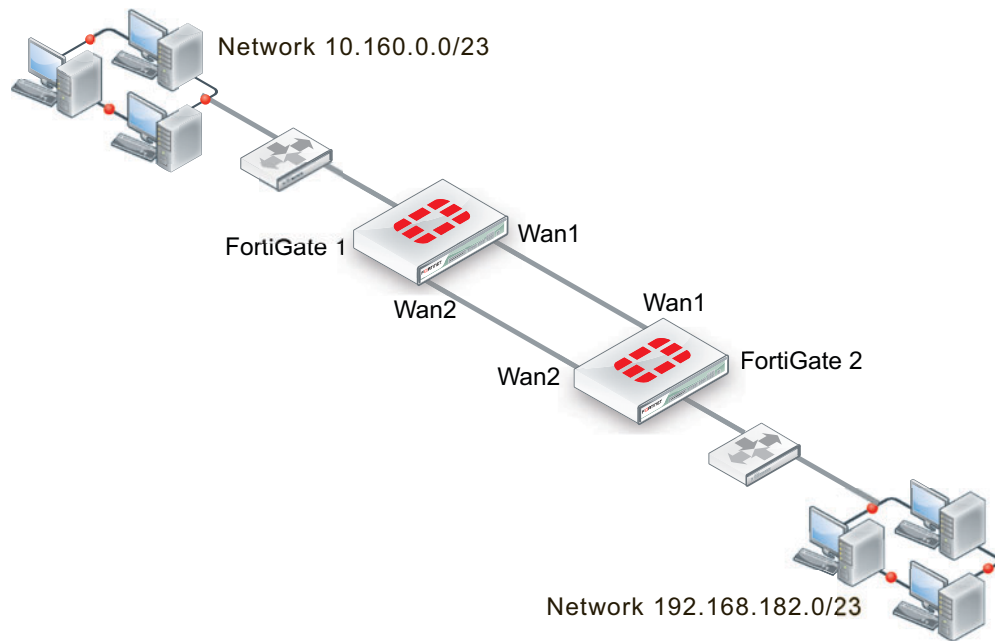
There are two main areas to test in this network configuration; network connectivity, and OSPF routing.

To test the network connectivity, see if computers on the Accounting or R&D networks can access the internet. If you need to troubleshoot network connectivity, see the FortiOS Handbook Troubleshooting chapter.

To test the OSPF routing, check the routing tables on the FortiGate units to ensure the expected OSPF routes are present. If you need help troubleshooting OSPF routing, see [“Troubleshooting OSPF” on page 148](#).

Controlling redundant links by cost

In this scenario, two FortiGate units have redundant links: one link between their WAN1 interfaces and another between their WAN2 interfaces.



FortiGate 1 should learn the route to network 192.168.182.0 and FortiGate 2 should learn the route to network 10.160.0.0. Under normal conditions, they should learn these routes through the WAN1 link. The WAN2 link should be used only as a backup.

With the default settings, each FortiGate unit learns these routes from both WAN1 and WAN2.

FortiGate 1:

```
FGT1 # get router info ospf neighbor
OSPF process 0:
Neighbor ID Pri State Dead Time Address Interface
10.2.2.2 1 Full/Backup 00:00:33 10.182.0.187 wan1
10.2.2.2 1 Full/Backup 00:00:31 10.183.0.187 wan2

FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.183.0.187, wan2, 00:00:01
[110/10] via 10.182.0.187, wan1, 00:00:01
O 192.168.182.0/23 [110/20] via 10.183.0.187, wan2, 00:02:04
[110/20] via 10.182.0.187, wan1, 00:02:04
```

FortiGate 2:

```
FGT2 # get router info ospf neighbor
OSPF process 0:
Neighbor ID Pri State Dead Time Address Interface
10.1.1.1 1 Full/DR 00:00:38 10.182.0.57 wan1
10.1.1.1 1 Full/DR 00:00:38 10.183.0.57 wan2

FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/20] via 10.183.0.57, wan2, 00:00:39
[110/20] via 10.182.0.57, wan1, 00:00:39
```

Adjusting the route costs

On both FortiGate units, the cost of the route through WAN2 is adjusted higher so that this route will only be used if the route through WAN1 is unavailable. The default cost is 10. The WAN2 route will be changed to a cost of 200.

On both FortiGate units:

```
config router ospf
  config ospf-interface
    edit "WAN2_higher_cost"
      set cost 200
      set interface "wan2"
    end
```

Now both FortiGate units use only the WAN1 route:

FortiGate 1:

```
FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.182.0.187, wan1, 00:00:40
O 192.168.182.0/23 [110/20] via 10.182.0.187, wan1, 00:00:40
```

FortiGate 2:

```
FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/20] via 10.182.0.57, wan1, 00:09:37
```

LSDB check on FortiGate 1:

```
FGT1 # get router info ospf database router lsa
Router Link States (Area 0.0.0.0)
LS age: 81
Options: 0x2 (*|-|-|-|-|E|-)
Flags: 0x0
LS Type: router-LSA
Link State ID: 10.1.1.1
Advertising Router: 10.1.1.1
LS Seq Number: 8000000b
Checksum: 0xe637
Length: 60
Number of Links: 3
```

```
Link connected to: Stub Network
(Link ID) Network/subnet number: 10.160.0.0
(Link Data) Network Mask: 255.255.254.0
Number of TOS metrics: 0
```

```
TOS 0 Metric: 10
Link connected to: a Transit Network
(Link ID) Designated Router address: 10.183.0.187
(Link Data) Router Interface address: 10.183.0.57
Number of TOS metrics: 0
TOS 0 Metric: 200
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 10.182.0.57
(Link Data) Router Interface address: 10.182.0.57
Number of TOS metrics: 0
TOS 0 Metric: 10
```

```
LS age: 83
Options: 0x2 (*|-|-|-|-|E|-)
Flags: 0x2 : ASBR
LS Type: router-LSA
Link State ID: 10.2.2.2
Advertising Router: 10.2.2.2
LS Seq Number: 8000000e
Checksum: 0xfc9b
Length: 60
Number of Links: 3
```

```
Link connected to: Stub Network
(Link ID) Network/subnet number: 192.168.182.0
(Link Data) Network Mask: 255.255.254.0
Number of TOS metrics: 0
TOS 0 Metric: 10
```



```
Link connected to: a Transit Network
(Link ID) Designated Router address: 10.183.0.187
(Link Data) Router Interface address: 10.183.0.187
Number of TOS metrics: 0
TOS 0 Metric: 200
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 10.182.0.57
(Link Data) Router Interface address: 10.182.0.187
Number of TOS metrics: 0
TOS 0 Metric: 10
```

Verifying route redundancy

Bring down WAN1 and then check the routes on the two FortiGate units.

FortiGate 1:

```
FGT1 # get router info routing-table ospf
FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.183.0.187, wan2, 00:00:06
O 192.168.182.0/23 [110/210] via 10.183.0.187, wan2, 00:00:06
```

FortiGate 2:

```
FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/210] via 10.183.0.57, wan2, 00:00:14
```

The WAN2 interface is now in use on both units.

Intermediate System to Intermediate System Protocol (IS-IS)

This section describes the Intermediate System to Intermediate System Protocol (IS-IS).

The following topics are included in this section:

- [IS-IS background and concepts](#)
- [Troubleshooting IS-IS](#)
- [Simple IS-IS example](#)

IS-IS background and concepts

Intermediate System to Intermediate System Protocol (IS-IS) allows routing of ISO's OSI protocol stack Connectionless Network Service (CLNS). IS-IS is an Interior Gateway Protocol (IGP) not intended to be used between Autonomous Systems (ASes).

This section contains:

- [Background](#)
- [How IS-IS works](#)
- [Parts and terminology of IS-IS](#)

Background

IS-IS was developed by Digital Equipment Corporation and later standardized by ISO in 1992 as ISO 19589 (see [RFC 1142](#) — note this RFC is different from the ISO version). At roughly the same time, the Internet Engineering Task Force developed OSPF (see "[Open Shortest Path First \(OSPF\)](#)" on [page 136](#)). After the initial version, IP support was added to IS-IS and this version was called Integrated IS-IS (see [RFC 1195](#)). Its widespread use started when an early version of IS-IS was included with BSD v4.3 Linux as the routed daemon. The routing algorithm used by IS-IS, the Bellman–Ford algorithm, first saw widespread use as the initial routing algorithm of the ARPANET.

IS-IS is a link state protocol well-suited to smaller networks that is in widespread use and has near universal support on routing hardware. It is quick to configure, and works well if there are no redundant paths. However, IS-IS updates are sent out node-by-node, so it can be slow to find a path around network outages. IS-IS also lacks good authentication, can not choose routes based on different quality of service methods, and can create network loops if you are not careful. IS-IS uses Dijkstra's algorithm to find the best path, like OSPF.

While OSPF is more widely known, IS-IS is a viable alternative to OSPF in enterprise networks and ISP infrastructures, largely due to its native support for IPv6 and its non-disruptive methods for splitting, merging, migrating, and renumbering network areas.

The FortiGate implementation supports both IS-IS (see [RFCs 1142](#) and [1162](#)) and Integrated IS-IS (see [RFCs 1195](#) and [5308](#)).

How IS-IS works

As one of the original modern dynamic routing protocols, IS-IS is straightforward. Its routing algorithm is not complex, there are some options to allow fine tuning, and it is straightforward to configure IS-IS on FortiGate units.

From RFC 1142:

The routing algorithm used by the Decision Process is a shortest path first (SPF) algorithm. Instances of the algorithm are run independently and concurrently by all intermediate systems in a routing domain. IntraDomain routing of a PDU occurs on a hop-by-hop basis: that is, the algorithm determines only the next hop, not the complete path, that a data PDU will take to reach its destination.

IS-IS versus static routing

IS-IS was one of the earliest dynamic routing protocols to work with IP addresses. As such, it is not as complex as more recent protocols. However, IS-IS is a big step forward from simple static routing.

While IS-IS may be slow in response to network outages, static routing has zero response. The same is true for convergence—static routing has zero convergence. Both IS-IS and static routing have the limited hop count, so it is neither a strength nor a weakness.

TLV

IS-IS uses *type-length-variable (TLV)* parameters to carry information in Link-State PDUs (LSPs). Each IS-IS LSP consists of a variable-length header to which TLVs are appended in order to extend IS-IS for IP routing. The TLV field consists of one octet of type (T), one octet of length (L), and “L” octets of Value (V). They are included in all of the IS-IS [Packet types](#). For a complete breakdown of the LSP, refer to “[LSP structure](#)” on [page 179](#).

In IS-IS, TLVs are used to determine route-leaking and authentication, and are also used for IPv4 and IPv6 awareness and reachability.

- To determine which TLVs are responsible for route-leaking, refer to “[Default routing](#)” on [page 182](#).
- To determine which TLVs are responsible for authentication, refer to “[Authentication](#)” on [page 183](#).
- To determine which TLVs are responsible for IPv4 and IPv6 awareness and reachability, refer to “[Integrated IS-IS](#)” on [page 184](#).

For a complete list of reserved TLV codepoints, refer to [RFC 3359](#).

LSP structure

It is difficult to fully understand a routing protocol without knowing what information is carried in its packets. Knowing how routers exchange each type of information will help you better understand the IS-IS protocol and will allow you to configure your network more appropriately.

This section provides information on the contents of the IS-IS LSP. LSPs describe the network topology and can include IP routes and checksums.

NSAP and NET

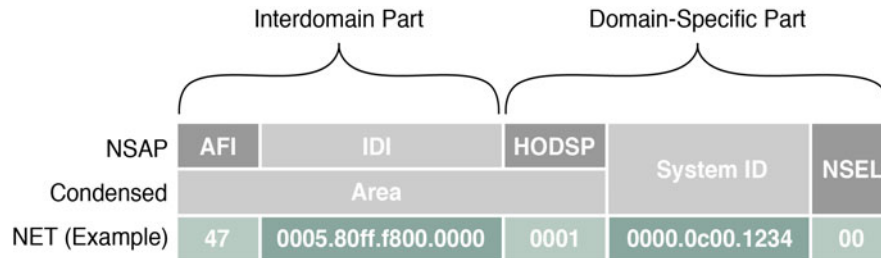
IS-IS routing protocol utilizes ISO network addressing to identify network interfaces. The addresses are known as Network Service Access Points (NSAPs). In general, IS-IS routers consist of only one NSAP, whereas IP addressing requires one IP address per interface.

In IS-IS, the NSAP address is translated into a Network Entity Title (NET), which is the same as the NSAP but can differentiate end systems by way of a byte called the *n-selector* (NSEL). In

order for adjacencies to form in IS-IS, the NSEL must necessarily be set to zero, to indicate “this system”. The total NET can be anywhere between 8 and 20 bytes long due to the support for variable length area addressing.

The following diagram identifies the individual parts of the NSAP, with explanations below.

Table 14: NSAP and NET example



AFI — The *Authority and Format Identifier (AFI)* specifies the format of the addressing family used. IS-IS is designed to carry routing information for several different protocols. Each entry has an address family identifier that identifies the globally unique Interdomain Part (IDP). For example, 49 is the AFI for private addresses, whereas 47 is the AFI for international organizations.

IDI — The *Initial Domain Identifier (IDI)* identifies the routing domain within an interconnected network. The length of the IDI is typically determined by the AFI. If you are using an AFI of 49, you do not need to specify an IDI, since the network is private.

HODSP — The *High Order Domain-Specific Part (HODSP)* identifies the unique address within a specific routing domain. Together, the AFI, IDI, and HODSP define the area address. All of the nodes within an area must have the same area address.

System ID — The *System ID* represents the 6-8 byte router identifier. The ID could be Media Access Control (MAC) format, as in the example above, or a static length IP address expressed in binary-coded decimal (BCD) format.

NSEL — The *n-selector (NSEL)*, as previously described, identifies the network layer transport service and must always be set to zero for IS-IS NETs.

Parts and terminology of IS-IS

Before you can understand how IS-IS functions, you need to understand some of the main concepts and parts of IS-IS.

This section includes:

- [DIS election and pseudonode LSP](#)
- [Packet types](#)
- [Default routing](#)
- [Timer options](#)
- [Authentication](#)
- [Integrated IS-IS](#)

DIS election and pseudonode LSP

In IS-IS routing protocol, a single router is chosen to be the designated intermediate system (DIS). The election of the DIS is determined automatically and dynamically on the LAN depending on highest interface priority and the subnetwork point of attachment (SNPA). The FortiGate is typically the DIS, and each router in its LAN is an intermediate system (IS).

Unlike OSPF, which elects a designated router (DR) and backup designated router (BDR), the DIS has no backup and determines the election of a new DIS whenever a router is added to the LAN or whenever the current DIS drops. A backup DIS is irrelevant since all of the routers on an IS-IS system are synchronized, and the short Hello interval used by the DIS quickly detects failures and the subsequent replacement of the DIS.

Synchronization of all the nodes in an IS-IS area could prove troublesome when updating the network infrastructure, and would demand ever-increasing resources each time a new router is added (at an exponential scale). For this purpose the DIS creates a pseudonode, which is essentially a virtual, logical node representing the LAN. The pseudonode requests adjacency status from all the routers in a multi-access network by sending IS-IS Hello (IIH) PDUs to Level 1 and Level 2 routers (where Level 1 routers share the same address as the DIS and Level 2 routers do not). Using a pseudonode to alter the representation of the LAN in the link-state database (LSD) greatly reduces the amount of adjacencies that area routers have to report. In essence, a pseudonode *collapses* a LAN topology, which allows a more linear scale to link-state advertising.

In order to maintain the database synchronization, the DIS periodically sends complete sequence number packets (CSNPs) to all participating routers.

Packet types

Four general packet types (PDUs) are communicated through IS-IS, appearing at both Level 1 and Level 2. They are described below.

Intermediate System-to-Intermediate System Hello (IIH) PDU — As mentioned previously, the IIH PDU, or Hello packet, detects neighbouring routers and indicates to the pseudonode the area's adjacency mesh. The Hello packet, flooded to the multicast address, contains the system ID of the sending router, the holding time, the circuit type of the interface on which the PDU was sent, the PDU length, the DIS identifier, and the interface priority (used in DIS election). The Hello packet also informs its area routers that it is the DIS.

Hello packets are padded to the maximum IS-IS PDU size of 1492 bytes (the full MTU size) to assist in the detection of transmission errors with large frames or with MTU mismatches between adjacencies.

The DIS typically floods Hello packets to the entire LAN every three seconds.

Link-state PDU (LSP) — The LSP contains information about each router in an area and its connected interfaces. LSPs are refreshed periodically and acknowledged on the network by way of sequence number PDUs. If new LSP information is found, based on the most recent complete sequence number PDU (CSNP), then out-of-date entries in the link-state database (LSDB) are removed and the LSDB is updated.

For a more detailed breakdown of the LSP, see [“LSP structure” on page 179](#).

Complete sequence number PDU (CSNP) — CSNPs contain a list of all LSPs in the current LSDB. The CSNP informs other area routers of missing or outdated links in the adjacency mesh. The receiving routers then use this information to update their own database to ensure that all area routers converge.

In contrast to Hello packets, CSNPs are sent every ten seconds and only between neighbours. In other words, they are never flooded.

Partial sequence number PDU (PSNP) — PSNPs are used to request and acknowledge LSP information from an adjacency. When a router compares a CSNP with its local database and determines a discrepancy, the router requests an updated LSP using a PSNP. Once received, the router stores the LSP in its local database and responds to the DIS with acknowledgement.

Default routing

The default route is used if either there are no other routes in the routing table or if none of the other routes apply to a destination. Including the gateway in the default route gives all traffic a next-hop address to use when leaving the local network. The gateway address is normally another router on the edge of the local network.

FortiGate units come with a default static route with an IPv4 address of 0.0.0.0, an administration distance of 10, and a gateway IPv4 address. Beginner administrators can use the default route settings until a more advanced configuration is warranted.

By default, all routes are displayed in the Routing Monitor list. To display the routes in the routing table, go to *Route > Monitor > Routing Monitor*.

Route leaking

Route leaking is the term used to describe the bi-directional flow of information between internal and external routing interfaces. By default, IS-IS leaks routing information from a Level 1 area into a Level 2 area. In order to leak Level 2 routing information into a Level 1 area, you must configure an export policy. Whether or not a route is leaked is determined by the ATT bit, using TLV 128 (for internal IP reachability) and TLV 130 (for external IP address information). For more information on TLVs, see “[Troubleshooting IS-IS](#)” on page 185.

To configure IS-IS route leaking, use the following CLI commands.

1. On a Level 1-2 router:

```
config router isis
    set redistribute-l2 enable
end
```

2. On a Level 1 router:

```
config router isis
    get router info routing-table isis
    get router info isis route
end
```

Default information originate option

Enabling default-information-originate generates and advertises a default route into the FortiGate unit’s IS-IS-enabled networks. The generated route may be based on routes learned through a dynamic routing protocol, routes in the routing table, or both. IS-IS does not create the default route unless you use the `always` option.

Select *Disable* if you experience any issues or if you wish to advertise your own static routes into IS-IS updates.

The CLI commands associated with default information originate include:

```
config router isis
    set default-originate
end
```

Timer options

IS-IS uses various timers to regulate its performance including a garbage timer, update timer, and timeout timer. The FortiGate unit default timer settings (30, 180, and 120 seconds respectively) are effective in most configurations—if you change these settings, ensure that the new settings are compatible with local routers and access servers.

You can configure the three IS-IS timers in the CLI, using the following commands:

```
config router isis
  set garbage-timer
  set update-timer
  set timeout-timer
end
```

You will find more information on each timer below.

Garbage timer

The garbage timer is the amount of time (in seconds) that the DIS will advertise a route as being unreachable before deleting the route from the routing table. If this timer is shorter, it will keep more up-to-date routes in the routing table and remove old ones faster. This results in a smaller routing table which is useful if you have a very large network, or if your network changes frequently.

Update timer

The update timer determines the interval between routing updates. Generally, this value is set to 30 seconds. There is some randomness added to help prevent network traffic congestion, which could result from all routers simultaneously attempting to update their neighbors. The update timer should be at least three times smaller than the timeout timer, otherwise you will experience an error.

If you are experiencing significant traffic on your network, you can increase this interval to send fewer updates per minute. However, ensure you increase the interval for all the routers on your network or you will experience timeouts that will degrade your network speed.

Timeout timer

The timeout timer is the maximum amount of time (in seconds) that a route is considered reachable while no updates are received for the route. This is the maximum time the DIS will keep a reachable route in the routing table while no updates for that route are received. If the DIS receives an update for the route before the timeout period expires, the timer is restarted. The timeout period should be at least three times longer than the update period, otherwise you will experience an error.

If you are experiencing problems with routers not responding in time to updates, increase this timer. However, remember that longer timeout intervals result in longer overall update periods — it may be considerable time before the DIS is done waiting for all the timers to expire on unresponsive routes.

Authentication

In routing protocols, it is typically desirable to establish authentication rules that prevent malicious and otherwise unwanted information from being injected into the routing table. IS-IS routing protocol utilizes TLV 10 to establish authentication. For more information on TLVs, see “TLV” on page 179.

Initially, IS-IS used plain Clear Text to navigate the authentication rules, but this was found to be insecure since the Clear Text packets were unencrypted and could be exposed to packet sniffers. As per [RFC 3567](#), HMAC-MD5 and Enhanced Clear Text authentication features were introduced to IS-IS, both of which encrypt authentication data, making them considerably more secure than using plain Clear Text authentication.

HMAC-MD5 authentication

Hashed Message Authentication Codes - Message Digest 5 (HMAC-MD5) is a mechanism for applying a cryptographic hash function to the message authentication process. It is applied at

both Level 1 and Level 2 routing. In IS-IS, an HMAC-MD5 can be applied to each type of LSP, on different interfaces, and with different passwords.

Authentication data is hashed using an AH (Authentication Header) key. From [RFC 2085](#):

The “AH Key” is used as a shared secret between two communicating parties. The Key is not a “cryptographic key” as used in a traditional sense. Instead, the AH key (shared secret) is hashed with the transmitted data and thus, assures that an intervening party cannot duplicate the authentication data. [...] Implementation should, and as frequently as possible, change the AH key. Keys need to be chosen at random, or generated using a cryptographically strong pseudo-random generator seeded with a random seed.”

Clear Text authentication uses the configuration commands `area-password` and `domain-password` for authentication, but when migrating from Clear Text authentication to HMAC-MD5, these command settings are automatically overwritten.

By the year 2005, the MD5 hash function had been identified as vulnerable to collision search attacks and various weaknesses. While such vulnerabilities do not compromise the use of MD5 within HMAC, administrators need to be aware of potential developments in cryptanalysis and cryptographic hash functions in the likely event that the underlying hash function needs to be replaced.

Enhanced Clear Text authentication

Enhanced Clear Text authentication is an extension to Clear Text authentication that allows the encryption of passwords as they are displayed in the configuration. It includes a series of authentication mode commands and an authentican key chain, and allows for more simple password modification and password management. Enhanced Clear Text authentication also provides for smoother migration to and from changing authentication types. Intermediate systems continue to use the original authentication method until all the area routers are updated to use the new method.

Authentication key chain

A key chain is a list of one or more authentication keys including the send and receive lifetimes for each key. Keys are used for authenticating routing packets only during the specified lifetimes. A router migrates from one key to the next according to the scheduled send and receive lifetimes. If an active key is unavailable, then the PDU is automatically discarded.

From [RFC 5310](#):

It should be noted that the cryptographic strength of the HMAC depends upon the cryptographic strength of the underlying hash function and on the size and quality of the key.

Integrated IS-IS

Integrated IS-IS is an extended version of IS-IS that includes support for both IPv4 and IPv6. IPv4 and IPv6 interface addresses are determined by TLVs 132 and 232, respectively. The parameter responsible for IPv6 reachability is TLV 236. For more information on TLVs, see “[Troubleshooting IS-IS](#)” on page 185.

The FortiGate unit command `config router isis` is almost the same except that IPv6 addresses are used. Also, if you are going to use prefix or access lists with Integrated IS-IS, you must use the `config router access-list6` or `config prefix-list6` versions of those commands.

If you want to route IPv4 traffic over an IPv6 network, you can use the command `config system ipv6-tunnel` to configure the FortiGate unit to do this. The IPv6 interface is configured under `config system interface`. All subnets between the source and destination addresses must support IPv6. This command is not supported in Transparent mode.

For example, you want to set up a tunnel on the port1 interface starting at 2002:C0A8:3201:: on your local network and tunnel it to address 2002:A0A:A01:: where it will need access to an IPv4 network again. Use the following command:

```
config system ipv6-tunnel
  edit test_tunnel
    set destination 2002:A0A:A01::
    set interface port1
    set source 2002:C0A8:3201::
  end
end
```

The CLI commands associated with Integrated IS-IS include:

```
config router isis
config router access-list6
config router prefix-list6
config system ipv6-tunnel
get router info6 *
```

Troubleshooting IS-IS

If you want to troubleshoot Integrated IS-IS, it is the same as with IS-IS but you must specify the different protocol, and use IPv6 addresses. This applies to commands such as `get router info6` when you want to see the routing table, or other related information.

This section includes:

- [Routing loops](#)
- [Split horizon and Poison reverse updates](#)

Routing loops

Normally in routing, a path between two addresses is chosen and traffic is routed along that path from one address to the other. When there is a routing loop, that normal path doubles back on itself creating a loop. When there are loops, the network has problems.

A routing loop happens when a normally functioning network has an outage, and one or more routers are offline. When packets encounter this, an alternate route is attempted to maneuver around the outage. During this phase it is possible for a route to be attempted that involves going back a hop, and trying a different hop forward. If that hop forward is blocked by the outage as well, a hop back and possibly the original hop forward may be selected. You can see if this continues, how it can consume not only network bandwidth but also many resources on those routers affected. The worst part is this situation will continue until the network administrator changes the router settings, or the downed routers come back online.

Routing loop effect on the network

In addition to this “traffic jam” of routed packets, every time the routing table for a router changes that router sends an update out to all of the IS-IS routers connected to it. In a network loop, its possible for a router to change its routes very quickly as it tries and fails along these new routes. This can quickly result in a flood of updates being sent out, which can effectively grind the network to a halt until the problem is fixed.

How can you spot a routing loop

Any time network traffic slows down, you will be asking yourself if it is a network loop or not. Often slowdowns are normal, they are not a full stoppage, and normal traffic resumes in a short period of time.

If the slow down is a full halt of traffic or a major slowdown does not return to normal quickly, you need to do serious troubleshooting quickly.

Some methods to troubleshoot your outage include:

- [Checking your logs](#)
- [Using SNMP network monitoring](#)
- [Using dead gateway detection and e-mail alerts](#)
- [Looking at the packet flow](#)

If you aren't running SNMP, dead gateway detection, or you have non-Fortinet routers in your network, you can use networking tools such as ping and traceroute to define the outage on your network and begin to fix it.

Checking your logs

If your routers log events to a central location, it can be easy to check the logs for your network for any outages.

On your FortiGate unit, go to *Log & Report > Log & Archive Access*. You will want to look at both event logs and traffic logs. Events to look for will generally fall under CPU and memory usage, interfaces going offline (due to dead gateway detection), and other similar system events.

Once you have found and fixed your network problem, you can go back to the logs and create a report to better see how things developed during the problem. This type of forensics analysis can better help you prepare for next time.

Using SNMP network monitoring

If your network had no problems one minute and slows to a halt the next, chances are something changed to cause that problem. Most of the time an offline router is the cause, and once you find that router and bring it back online, things will return to normal.

If you can enable a hardware monitoring system such as SNMP or sFlow on your routers, you can be notified of the outage and where it is exactly as soon as it happens.

Ideally you can configure SNMP on all your FortiGate routers and be alerted to all outages as they occur.

To use SNMP to detect potential routing loops

1. Go to *System > Config > SNMP*.
2. Enable *SNMP Agent*.
3. Optionally enter the *Description*, *Location*, and *Contact* information for this device for easier location of the problem report.
4. In either *SNMP v1/v2c* section or *SNMP v3* section, as appropriate, select *Create New*.
5. Enter the *Community Name* that you want to use.
6. In *Hosts*, select *Add* to add an IP address where you will be monitoring the FortiGate unit. You can add up to 8 different addresses.
7. Ensure that ports 161 and 162 (SNMP queries and traps) are allowed through your security policies.
8. In *SNMP Event*, select the events you want to be notified of. For routing loops this should include *CPU Overusage*, *Memory Low*, and possibly *Log disk space low*. If there are

problems, the log will be filling up quickly, and the FortiGate unit's resources will be overused.

9. Select *OK*.
10. Configure SNMP host (manager) software on your administration computer. This will monitor the SNMP information sent out by the FortiGate unit. Typically you can configure this software to alert you to outages or CPU spikes that may indicate a routing loop.

Using dead gateway detection and e-mail alerts

Another tool available to you on FortiGate units is the dead gateway detection. This feature allows the FortiGate unit to ping a gateway at regular intervals to ensure it is online and working. When the gateway is not accessible, that interface is marked as down.

To detect possible routing loops with dead gateway detection and e-mail alerts

1. To configure dead gateway detection, go to *Router > Static > Settings* and select *Create New*.
2. Set the *Ping Interval* (how often to send a ping), and *Failover Threshold* (how many lost pings is considered a failure). A smaller interval and smaller number of lost pings will result in faster detection, but will create more traffic on your network.
3. To configure interface status change notification, go to *Log & Report > Log Config > Alert E-mail*.
4. After you enter your email details, select the events you want to be alerted about — in our case *Configuration changes*. You may also want to log CPU and Memory usage as a network outage will cause your CPU activity to spike.



If you have VDOMs configured, you will have to enter the basic SMTP server information in the Global section, and the rest of the configuration within the VDOM that includes this interface.

After this configuration, when this interface on the FortiGate unit cannot connect to the next router, the FortiGate unit will bring down the interface and alert you with an email to the outage.

Looking at the packet flow

If you want to see what is happening on your network, look at the packets travelling on the network. In this situation, you are looking for routes that have metrics higher than 15 as that indicates they are unreachable. Ideally if you debug the flow of the packets, and record the routes that are unreachable, you can create an accurate picture of the network outage.

Action to take on discovering a routing loop

Once you have mapped the problem on your network, and determined it is in fact a routing loop there are a number of steps to take in correcting it.

1. Get any offline routers back online. This may be a simple reboot, or you may have to replace hardware. Often this first step will restore your network to its normal operation, once the routing tables finish being updated.
2. Change your routing configuration on the edges of the outage. Even if step 1 brought your network back online, you should consider making changes to improve your network before the next outage occurs. These changes can include configuring features like hold downs and triggers for updates, split horizon, and poison reverse updates.

Split horizon and Poison reverse updates

Split horizon is best explained with an example. You have three routers linked serially, let's call them A, B, and C. A is only linked to B, C is only linked to B, and B is linked to both A and C. To get to C, A must go through B. If the link to C goes down, it is possible that B will try to use A's route to get to C. This route is A-B-C, so it will not work. However, if B tries to use it this begins an endless loop.

This situation is called a split horizon because from B's point of view the horizon stretches out in each direction, but in reality it only is on one side.

Poison reverse is the method used to prevent routes from running into split horizon problems. Poison reverse "poisons" routes away from the destination that use the current router in their route to the destination. This "poisoned" route is marked as unreachable for routers that cannot use it. In IS-IS this means that route is marked with a distance of 16.

Simple IS-IS example

This is an example of a typical medium-sized network configuration using IS-IS routing.

Imagine a company with four FortiGate devices connected to one another. A FortiGate at one end of the network connects to two routers, each with its own local subnet; one of these routers uses OSPF and the other uses RIP.

Your task is to configure the four FortiGates to route traffic and process network updates using IS-IS, such that the farthest FortiGate (see 'FGT4' in [Network layout and assumptions](#)) receives route updates for the two routers at the opposite end of the network. Furthermore, FGT4 has been given a loopback subnet that must be identified by the router running RIP.

Since the internal networks use OSPF and RIP, those protocols will need to be redistributed through the IS-IS network. To keep the example simple, there will be no authentication of router traffic.

With IS-IS properly configured in this example, if a router fails or temporarily goes offline, the route change will propagate throughout the system.

This section includes the following topics:

- [Network layout and assumptions](#)
- [Expectations](#)
- [CLI configuration](#)
- [Verification](#)
- [Troubleshooting](#)

Network layout and assumptions

- It is assumed that each FortiGate is operating in NAT mode, running FortiOS 4.0MR2+.
- All interfaces have been previously assigned and no static routes are required.
- The AFI (Authority and Format Identifier) used is 49 : Locally administered (private).
- The Area identifiers are 0048 and 0049.

Expectations

- FGT4 must get the IS-IS route updates for RTR1 and RTR2 local subnets (10.1.1.0, 10.2.2.0, 10.3.3.0, 10.4.4.0).
- RTR1 must receive (via RIP2) the loopback subnet of FGT4 (10.60.60.1/32).

CLI configuration

The following CLI configuration occurs on each FortiGate (as identified), including only the relevant parts.

FGT1

```
config router isis
  config isis-interface
    edit "port3"
      set circuit-type level-1
      set network-type broadcast
      set status enable
    next
  end
config isis-net
  edit 1
    set net 49.0048.1921.6818.2136.00
  next
end
config redistribute "connected"
end
config redistribute "rip"
  set status enable
  set level level-1
end
config redistribute "ospf"
  set status enable
  set level level-1
end
end
config router rip
  config interface
    edit "port2"
      set receive-version 2
      set send-version 2
    next
  end
config network
  edit 1
    set prefix 10.10.10.0 255.255.255.0
  next
end
config redistribute "isis"
  set status enable
end
end
```

FGT2

```
config router isis
  config isis-interface
    edit "port3"
      set circuit-type level-1
      set network-type broadcast
      set status enable
    next
    edit "port2"
      set network-type broadcast
      set status enable
    next
  end
  config isis-net
    edit 1
      set net 49.0048.1221.6818.2110.00
    next
  end
  set redistribute-l1 enable
  set redistribute-l2 enable
end
```

FGT3

```
config router isis
  set is-type level-2-only
  config isis-interface
    edit "wan1"
      set network-type broadcast
      set status enable
    next
    edit "dmz1"
      set network-type broadcast
      set status enable
    next
  end
  config isis-net
    edit 1
      set net 49.0048.1921.6818.2108.00
    next
    edit 2
      set net 49.0049.1921.6818.2108.00
    next
  end
end
```

FGT4

```
config router isis
  set is-type level-2-only
  config isis-interface
    edit "wan1"
      set network-type broadcast
      set status enable
    next
  end
config isis-net
  edit 1
    set net 49.0049.1721.0160.1004.00
  next
end
config redistribute "connected"
  set status enable
end
end
```

Verification

Once the network has been configured, you need to test that it works as expected. Use the following CLI commands on the devices indicated.

Verifying if RTR1 receives loopback subnet of FGT4

```
(RTR1) # get router info routing-table all
```

Result:

```
C      10.1.1.0/24 is directly connected, vlan1
C      10.2.2.0/24 is directly connected, vlan2
C      10.10.10.0/24 is directly connected, dmz1
R      10.40.40.0/24 [120/2] via 10.10.10.1, dmz1, 00:04:07
R      10.50.50.0/24 [120/2] via 10.10.10.1, dmz1, 00:04:07
R      10.60.60.1/32 [120/2] via 10.10.10.1, dmz1, 00:04:07
```

(*) If required, filtering out 10.50.50.0 and 10.40.40.0 from the routing table could be done with a route-map.

Verification on FGT2, which is the border between L1 and L2 routing levels; looking at IS-IS information

```
FGT2 # get router info isis interface
```

Result:

```
port2 is up, line protocol is up
  Routing Protocol: IS-IS ((null))
  Network Type: Broadcast
  Circuit Type: level-1-2
  Local circuit ID: 0x01
  Extended Local circuit ID: 0x00000003
  Local SNPA: 0009.0f85.ad8c
```



```

IP interface address:
  10.40.40.2/24
IPv6 interface address:
Level-1 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.01
Number of active level-1 adjacencies: 0
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.01
Number of active level-2 adjacencies: 1
Next IS-IS LAN Level-1 Hello in 6 seconds
Next IS-IS LAN Level-2 Hello in 1 seconds
port3 is up, line protocol is up
  Routing Protocol: IS-IS ((null))
  Network Type: Broadcast
  Circuit Type: level-1
  Local circuit ID: 0x02
  Extended Local circuit ID: 0x00000004
  Local SNPA: 0009.0f85.ad8d
  IP interface address:
    10.30.30.2/24
  IPv6 interface address:
  Level-1 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.02
  Number of active level-1 adjacencies: 1
  Next IS-IS LAN Level-1 Hello in 2 seconds

```

```
FGT2 # get router info isis neighbour
```

Result:

System Id	Interface	SNPA	State	Holdtime	Type	Protocol
1921.6818.2108	port2	0009.0f04.0794	Up	22	L2	IS-IS
1921.6818.2136	port3	0009.0f85.acf7	Up	29	L1	IS-IS

Verification on FGT3, which is border between 2 areas; looking at IS-IS information

IS-IS router CLI commands available:

```
FGT3 # get router info isis ?
```

Result:

```

interface      show isis interfaces
neighbor       show CLNS neighbor adjacencies
is-neighbor    show IS neighbor adjacencies
database       show IS-IS link state database
route          show IS-IS IP routing table
topology       show IS-IS paths

```

Example of interface status and neighbours:

```
FGT3 # get router info isis interface
```

Result:

```

wan1 is up, line protocol is up
  Routing Protocol: IS-IS ((null))
  Network Type: Broadcast
  Circuit Type: level-1-2
  Local circuit ID: 0x01
  Extended Local circuit ID: 0x00000003

```

Local SNPA: 0009.0f04.0794
IP interface address:
10.40.40.1/24
IPv6 interface address:
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.01
Number of active level-2 adjacencies: 1
Next IS-IS LAN Level-2 Hello in 3 seconds

dmz1 is up, line protocol is up
Routing Protocol: IS-IS ((null))
Network Type: Broadcast
Circuit Type: level-1-2
Local circuit ID: 0x02
Extended Local circuit ID: 0x00000005
Local SNPA: 0009.0f04.0792
IP interface address:
10.50.50.1/24
IPv6 interface address:
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 1721.0160.1004.01
Number of active level-2 adjacencies: 1
Next IS-IS LAN Level-2 Hello in 7 seconds

FGT3 # **get router info isis neighbour**

Result:

System Id	Interface	SNPA	State	Holdtime	Type	Protocol
1221.6818.2110	wan1	0009.0f85.ad8c	Up	8	L2	IS-IS
1721.0160.1004	dmz1	0009.0f52.7704	Up	8	L2	IS-IS

Verification on FGT4 that the remote subnets from RTR1 and RTR2 are in the routing table and learned with IS-IS

FGT4 # **get router info routing-table all**

Result:

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default

```
i L2 10.1.1.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
i L2 10.2.2.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
i L2 10.3.3.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
i L2 10.4.4.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
i L2 10.10.10.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
i L2 10.11.11.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
i L2 10.20.20.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
i L2 10.30.30.0/24 [115/30] via 10.50.50.1, wan1, 00:13:55
i L2 10.40.40.0/24 [115/20] via 10.50.50.1, wan1, 00:15:30
C 10.50.50.0/24 is directly connected, wan1
C 10.60.60.1/32 is directly connected, loopback
```

Troubleshooting

The following diagnose commands are available for further IS-IS troubleshooting and will display all IS-IS activity (sent and received packets):

```
FGT # diagnose ip router isis level info
FGT # diagnose ip router isis all enable
FGT # diagnose debug enable
```

...to stop the debug type output :

```
FGT # diagnose ip router isis level none
```

Output and interpretation depends on the issue faced. You can provide this information to TAC if you open a support ticket.

Debugging IPv6 on IS-ISng

The debug command is very useful to see what is happening on the network at the packet level. The following CLI commands specify both IPv6 and IS-IS, so only IS-ISng packets will be reported. The output from these commands will show you the IS-ISng traffic on your FortiGate unit including RECV, SEND, and UPDATE actions.

```
FGT # diagnose ipv6 router isis level info
FGT # diagnose ipv6 router isis all enable
FGT # diagnose debug enable
```

These three commands will:

- turn on debugging in general
- set the debug level to information, a verbose reporting level
- turn on all IS-IS router settings

Part of the information displayed from the debugging is the metric (hop count). If the metric is 16, then that destination is unreachable since the maximum hop count is 15.

In general, you should see an update announcement, followed by the routing table being sent out, and a received reply in response.

Index

A

- access control list (ACL) 17
- administrative distance 19, 20
- anti-spoofing 19
- area 143
- AS
 - multihomed 54
 - number (ASN) 54
 - stub 54

B

- bgp
 - attribute
 - AS_PATH 105
 - ATOMIC_AGGREGATE 107
 - COMMUNITY 106
 - MULTI_EXIT_DESC 106
 - NEXT_HOP 106
 - BGP-4+ 98
 - clearing routes 100, 111
 - control plane 113
 - flap 113
 - graceful restart 113
 - MED 106
 - neighbors 100
 - password, MD5 100
 - RFC 1997 106
 - route reflectors (RR) 102
 - stabilizing the network 113
- Bi-directional Forwarding Detection (BFD) 114
- blackhole route 18

C

- Classless Inter-Domain Routing (CIDR) 45
- control plane 113
- convergence 44, 114

D

- dampening 112
 - reachability half-life 112
- dead gateway detection 72, 187
- defaultinfooriginate 182
- Dijkstra's algorithm 146
- distance vector protocols 46

E

- ECMP 18
- enhanced packet-matching 56
- equal cost multipath (ECMP) 18, 21, 37
- Exterior Gateway Protocol (EGP) 98
- exterior gateway protocol (EGP) 46

G

- Garbagetimer 182
- graceful restart 113

H

- HA 111
 - router monitor 12
 - routes 12

I

- IntegratedISIS 180
- interface
 - loopback 19
- interior gateway protocol (IGP) 46
- Internet Assigned Numbers Authority (IANA) 54
- IP, protocol 89 143
- ISISProtocol 178

L

- link-state advertisement (LSA) 46
- loopback interface 19

M

- Martian addresses 19
- missing MED 106
- Multi-Exit Discriminator (MED) 106
- Multipath routing 19

N

- not-so-stubby area (NSSA) 13

O

- ospf
 - adjacent routers 144, 149
 - area 143
 - area border router (ABR) 143
 - Dijkstra's algorithm 146
 - e1 13
 - e2 13
 - Hello packets 144
 - Hello protocol 144
 - IP datagrams 143
 - link-state 143
 - neighbor 144
 - NSSA 13
 - path cost 146
 - state of neighbor 149
- ospf AS 139

P

- Packettypes 181

- policy route
 - adding 10, 25
 - moving in list 27
- port 179 108
- protocol
 - ospf Hello 144
- Pseudonodes 180

R

- redistributed routes
 - ospf e1/e2 13
- reverse path lookup 19
- RFC
 - RFC 1349 26
 - RFC 1519 45
 - RFC 1771 98
 - RFC 1965 103
 - RFC 1966 102
 - RFC 1997 106
 - RFC 2385 100
 - RFC 2453 59
 - RFC 3065 103
 - RFC 3509 146
 - RFC 4271 98
 - RFC 4632 45
 - RFC 5237 25
 - RFC 791 26
- RIP
 - hop count 66
 - RFC 1058 59, 178
 - RFC 2453 59
 - RIP Next Generation (RIPng) 60, 184
 - version 1 59, 178
 - version 2 59
- route flap 113
 - HA 111
- route reflectors (RR) 102
- Routeleaking 182
- router monitor
 - HA 12

- routing
 - administrative distance 19
 - blackhole 18
 - domain 54
 - ECMP 18
 - enhanced packet-matching 56
 - loopback interface 19
 - routing table, searching 16
 - viewing information 11
- routing policy
 - protocol number 25
- routing table 146
 - removing routes 100

S

- Shortest Path First (SPF) 146
- Spill-over 40
- static route
 - adding policy 10, 25
 - administrative distance 19
 - moving in list 27
 - table priority 20
 - table sequence 20
- supernetting 104

T

- Timeouttimer 183
- troubleshooting
 - BFD 114
 - bgp 110
 - dampening 112
 - graceful restart 113
 - holddown timer 111
 - route flap 110
 - routing table 22
- Type of service (TOS) 26

U

- unicast reverse path forwarding (uRPF) 11
- Updatetimers 183
- usage-based ECMP 40