



FortiGate 6000F Administration

FortiOS 7.6.2



FORTINET DOCUMENT LIBRARY

<https://docs.fortinet.com>

FORTINET VIDEO LIBRARY

<https://video.fortinet.com>

FORTINET BLOG

<https://blog.fortinet.com>

CUSTOMER SERVICE & SUPPORT

<https://support.fortinet.com>

FORTINET TRAINING & CERTIFICATION PROGRAM

<https://www.fortinet.com/training-certification>

FORTINET TRAINING INSTITUTE

<https://training.fortinet.com>

FORTIGUARD LABS

<https://www.fortiguard.com>

END USER LICENSE AGREEMENT

<https://www.fortinet.com/doc/legal/EULA.pdf>

FEEDBACK

Email: techdoc@fortinet.com



January 28, 2025

FortiOS 7.6.2 FortiGate 6000F Administration Guide

01-762-465651-20250128

TABLE OF CONTENTS

Change log	8
What's New	9
What's new for FortiGate 6000F 7.6.2	9
What's new for FortiGate 6000F 7.6.1	9
What's new for FortiGate 6000F 7.6.0	9
FortiGate-6000 overview	10
Front panel interfaces	11
FortiGate-6000 schematic	11
Interface groups and changing data interface speeds	13
FortiGate 6000F series hardware generations	14
FortiGate 6001F model licensing	14
Applying your FortiGate 6001F FPC subscription license	14
Activating your FortiGate 6001F FPC perpetual license	15
Verifying your FortiGate 6001F FPC subscription and perpetual licenses	15
What to expect when your subscription license expires	15
Getting started with FortiGate 6000F series	17
Confirming startup status	18
FortiGate 6000F and the Security Fabric	19
Configuration synchronization	19
Confirming that FortiGate-6000 components are synchronized	20
Viewing more details about FortiGate-6000 synchronization	21
Cluster Status dashboard widget	22
FortiGate 6000F dashboard widgets	23
Cluster status	23
Resource Usage	23
Sensor Information	23
Multi VDOM mode	23
FortiGate-6000 7.6.2 incompatibilities and limitations	24
Remote console limitations	24
Default management VDOM	24
Maximum number of LAGs and interfaces per LAG	24
High Availability	24
FortiOS features that are not supported by FortiGate-6000 7.6.2	25
IPsec VPN tunnels terminated by the FortiGate 6000F	26
Traffic shaping and DDoS policies	26
FortiGuard web filtering and spam filtering queries	26
Web filtering quotas	26
Special notice for new deployment connectivity testing	26
Display the process name associated with a process ID	26
Managing individual FortiGate 6000F management boards and FPCs	27
Special management port numbers	27
HA mode special management port numbers	28
Connecting to individual FPC consoles	29

Connecting to individual FPC CLIs	30
Connecting to individual FPC CLIs of the secondary FortiGate-6000 in an HA configuration	30
Performing other operations on individual FPCs	30
Load balancing and flow rules	31
Setting the load balancing method	31
Determining the primary FPC	32
Flow rules for sessions that cannot be load balanced	32
Default flow rules for traffic that cannot be load balanced	33
GTP load balancing	34
Enabling GTP load balancing	34
GTP with dynamic source port allocation	35
GTP load balancing and fabric channel usage	36
PFCP load balancing	36
ICMP load balancing	37
Load balancing TCP, UDP, and ICMP sessions with fragmented packets	38
Adding flow rules to support DHCP relay	38
Flow rules to support multihop BFD (MBFD)	40
Flow rules to support IP multicast	41
Controlling SNAT port partitioning behavior	41
Showing how the DP3 processor will load balance a session	42
Maximum number of flow rules limited by hardware	43
SSL VPN load balancing	44
Setting up SSL VPN using flow rules	44
If you change the SSL VPN server listening port	45
Adding the SSL VPN server IP address	46
IPsec VPN load balancing	47
Configuring IPsec VPN load balancing	47
SD-WAN with multiple IPsec VPN tunnels	48
Example FortiGate 6000F IPsec VPN VRF configuration	48
Troubleshooting	49
FortiGate-6000 high availability	52
Introduction to FortiGate-6000 FGCP HA	52
Before you begin configuring HA	53
Connect the HA1 and HA2 interfaces for HA heartbeat communication	54
Default HA heartbeat VLAN triple-tagging	54
HA heartbeat VLAN double-tagging	56
Basic FortiGate-6000 HA configuration	58
Verifying that the cluster is operating normally	60
Confirming that the FortiGate-6000 HA cluster is synchronized	61
Viewing more details about HA cluster synchronization	62
Primary FortiGate-6000 selection with override disabled (default)	63
Primary FortiGate-6000 selection with override enabled	63
Failover protection	64
Device failure	64

Link failure	65
FPC failure	65
SSD failure	66
Session failover	67
Primary FortiGate-6000 recovery	67
HA reserved management interfaces	67
HA in-band management for management interfaces	68
Virtual clustering	68
Limitations of FortiGate-6000 virtual clustering	69
Virtual clustering VLAN/VDOM limitation	69
Configuring virtual clustering	70
HA cluster firmware upgrades	74
Distributed clustering	74
Modifying heartbeat timing	75
Changing the lost heartbeat threshold	76
Adjusting the heartbeat interval and lost heartbeat threshold	76
Changing the time to wait in the hello state	77
Changing how long routes stay in a cluster unit routing table	77
Session failover (session-pickup)	78
Enabling session pickup for TCP SCTP and connectionless sessions	78
If session pickup is disabled	79
Reducing the number of sessions that are synchronized	79
FortiGate-6000 FGSP	79
FGSP session synchronization options	80
Using data interfaces for FGSP session synchronization	81
Synchronizing sessions between FGCP clusters	82
Example FortiGate-6000 FGSP session synchronization with a data interface LAG	82
Example FortiGate-6000 FGSP session synchronization using HA interfaces	85
Standalone configuration synchronization	87
Selecting the config sync primary FortiGate-6000	88
Settings that are not synchronized	88
Limitations	89
FortiGate-6000 VRRP HA	89
Operating a FortiGate-6000	90
FortiLink support	90
ECMP support	90
VDOM-based session tables	90
IPv4 and IPv6 ECMP load balancing	91
Enabling auxiliary session support	91
ICAP support	91
Example ICAP configuration	92
SSL mirroring support	93
VXLAN support	94
FortiGate 6000F IPsec load balancing EMAC VLAN interface limitation	94
Configuring a FortiGate-6000 to operate in FIPS-CC mode	94
Troubleshooting FortiGate 6000F self tests	96
Global option for proxy-based certificate queries	98

Using data interfaces for management traffic	98
In-band management limitations	98
FortiGate 6000F management interface LAG and VLAN support	99
Management interface LAG limitations	99
Setting the MTU for a data interface	99
More management connections than expected for one device	100
More ARP queries than expected for one device - potential issue on large WiFi networks	100
VLAN ID 1 is reserved	100
Connecting to FPC CLIs using the console port	101
Firmware upgrade basics	101
Installing firmware on an individual FPC	102
Installing firmware from the BIOS after a reboot	104
Synchronizing the FPCs with the management board	105
FPC failover in a standalone FortiGate-6000	106
Troubleshooting an FPC failure	108
Displaying FPC link and heartbeat status	108
If both the base and fabric links are down	109
If only one link is down	109
Updating FPC firmware to match the management board	110
Troubleshooting configuration synchronization issues	110
Adjusting global DP3 timers	111
Changing the FortiGate-6301F and 6501F log disk and RAID configuration	111
Restarting the FortiGate 6000F	112
Packet sniffing for FPC and management board packets	112
Using the diagnose sniffer options slot command	113
Filtering out internal management traffic	113
Packet sniffing integrated switch fabric (ISF) interfaces	113
NMI switch and NMI reset commands	114
Diagnose debug flow trace for FPC and management board activity	114
FortiGate-6000 config CLI commands	116
config load-balance flow-rule	116
Syntax	116
config load-balance setting	119
config system console-server	123
Syntax	123
FortiGate-6000 execute CLI commands	124
execute factoryreset-shutdown	124
execute ha manage <id>	124
execute load-balance slot manage {<chassis>.slot slot}	124
execute load-balance slot nmi-reset <slot-map>	124
execute load-balance slot power-off <slot-map>	125
execute load-balance slot power-on <slot-map>	125
execute load-balance slot reboot <slot-map>	125
execute load-balance slot set-primary-worker <slot>	125
execute set-next-reboot rollback	125

execute system console-server {clearline connect showline}	126
execute system console-server clearline <line>	126
execute system console-server connect <slot>	126
execute system console-server showline	126
execute upload image {ftp tftp usb}	127

Change log

Date	Change description
January 28, 2025	FortiOS 7.6.2 document release.
November 28, 2024	FortiOS 7.6.1 document release.
July 25, 2024	FortiOS 7.6.0 document release.

What's New

This section describes what's been added to FortiOS 7.6 FortiGate 6000F releases.

What's new for FortiGate 6000F 7.6.2

FortiGate 6000F for FortiOS 7.6.2 includes the special notices and resolved issues described in the [FortiOS 7.6.2 Release Notes](#).

What's new for FortiGate 6000F 7.6.1

FortiGate 6000F for FortiOS 7.6.1 includes the resolved issues described in the [FortiOS 7.6.1 Release Notes](#).

What's new for FortiGate 6000F 7.6.0

FortiGate 6000F for FortiOS 7.6.0 includes the resolved issues described in the [FortiOS 7.6.0 Release Notes](#).

FortiGate-6000 overview

The FortiGate-6000 series is a collection of 3U 19-inch rackmount appliances that all include twenty-four 25GigE SFP28 and four 100GigE QSFP28 data network interfaces, as well as NP6 and CP9 processors to deliver high IPS/threat prevention performance.

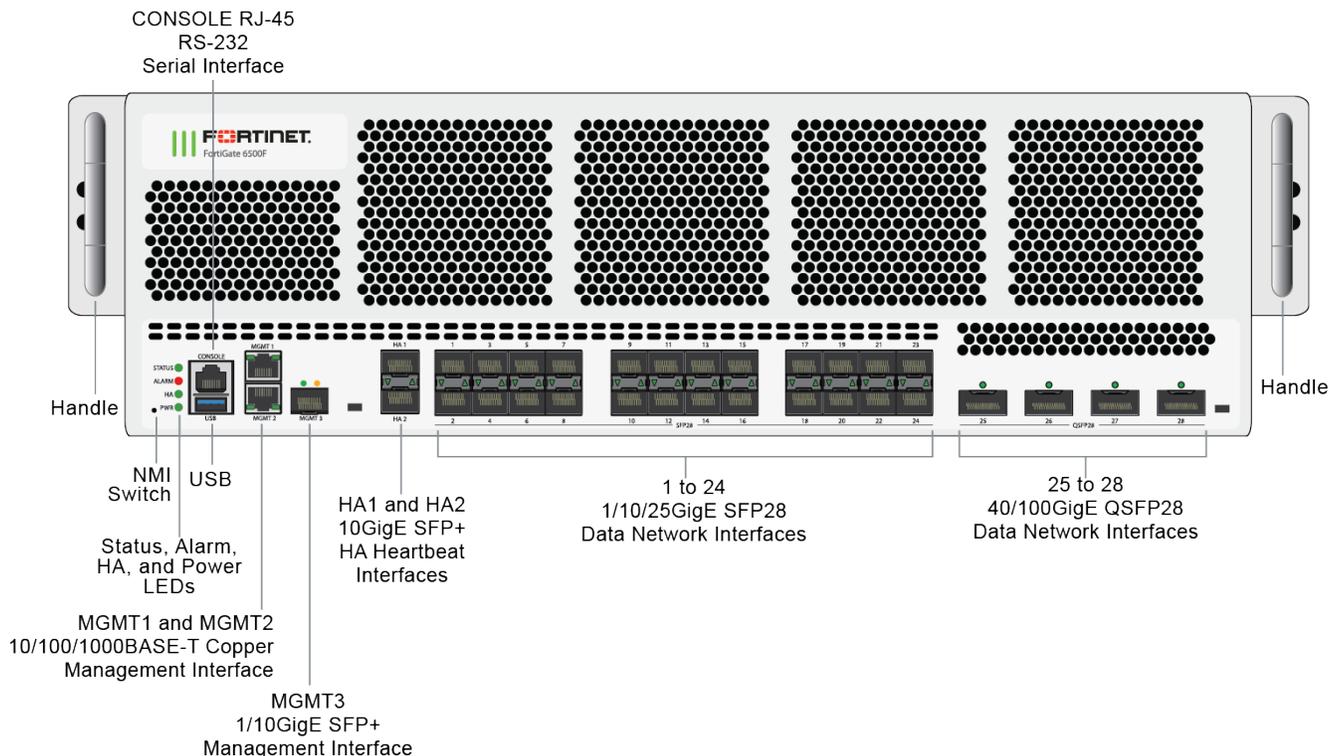
Currently, the following FortiGate-6000 models are available:

- FortiGate 6500F and FortiGate 6500F-DC
- FortiGate 6501F and FortiGate 6501F-DC
- FortiGate 6300F and FortiGate 6300F-DC
- FortiGate 6301F and FortiGate 6301F DC
- FortiGate 6001F and FortiGate 6001F-DC

All FortiGate 6000F series models have the same front and back panel configuration including the same network interfaces. The differences are the processing capacity of the individual models. All FortiGate 6000F series models include a management board (MBD) and internal Fortinet Processor Cards (FPCs) that contain NP6 and CP9 security processors. The management board handles management tasks, separating management tasks from data processing tasks that are handled by the FPCs. The FortiGate 6000F series uses session-aware load balancing to distribute sessions to the FPCs. The FortiGate-6500F includes ten FPCs and the FortiGate-6300F includes six FPCs.

Also the FortiGate 6001F, FortiGate 6501F, or FortiGate 6301F models include two internal 1 TByte log disks in a RAID-1 configuration.

FortiGate-6000 front panel (FortiGate-6500F shown)



Front panel interfaces

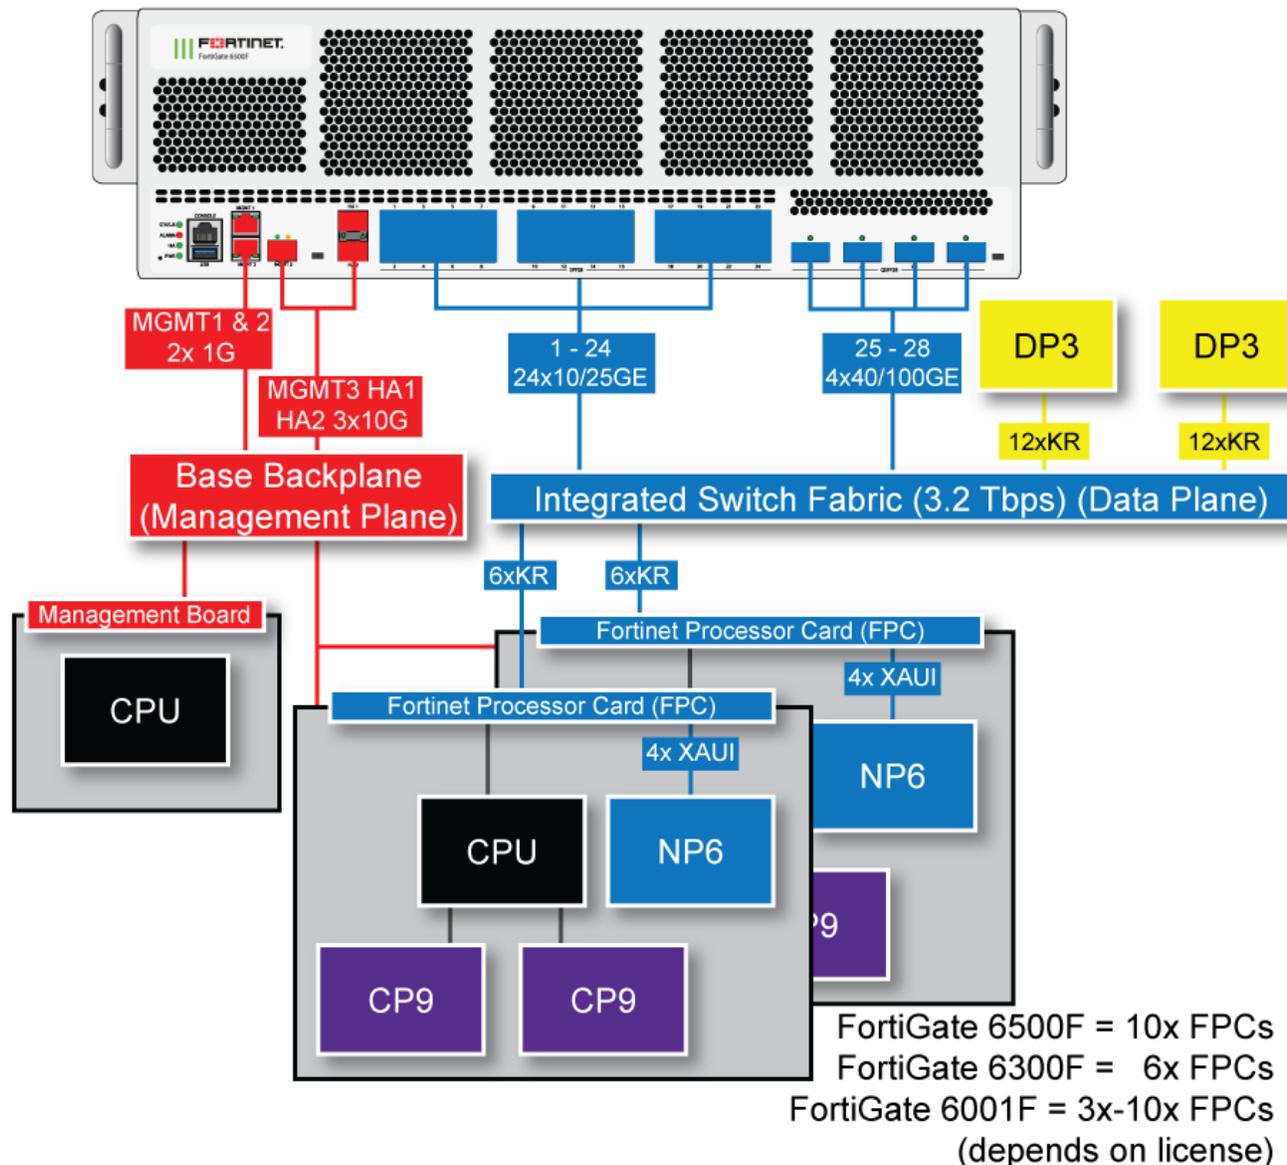
All FortiGate-6000 models have the following front panel interfaces:

- Twenty-four 1/10/25GigE SFP28 data network interfaces (1 to 24). The default speed of these interfaces is 10Gbps. These interfaces are divided into the following interface groups: 1 - 4, 5 - 8, 9 - 12, 13 - 16, 17 - 20, and 21 - 24. For information about interface groups, see [Interface groups and changing data interface speeds on page 13](#).
- Four 40/100GigE QSFP28 data network interfaces (25 to 28). The default speed of these interfaces is 40Gbps.
- Two front panel 1/10GigE SFP+ HA interfaces (HA1 and HA2) used for heartbeat, session sync, and management communication between two and only two FortiGate-6000s in an HA cluster. The default speed of these interfaces is 10Gbps. Operating them at 1Gbps is not recommended. A FortiGate-6000 cluster consists of two (and only two) FortiGate-6000s of the same model. To set up HA, you can use a direct cable connection between the FortiGate-6000s HA1 interfaces and a second direct cable connection between the HA2 interfaces. For more information about FortiGate-6000 HA, see [FortiGate-6000 high availability on page 52](#).
- Two 10/100/1000BASE-T out of band management Ethernet interfaces (MGMT1 and MGMT2).
- One front panel 1/10GigE SFP+ out of band management interface (MGMT3). You can use the 10Gbps MGMT3 interface for additional bandwidth management traffic that might be useful for high bandwidth activities such as remote logging.
- One RJ-45 RS-232 serial console connection.
- One USB connector.

FortiGate-6000 schematic

The FortiGate-6000 has separate data and management planes. The data plane handles all traffic and security processing functionality. The management plane handles management functions such as administrator logins, configuration and session synchronization, SNMP and other monitoring, HA heartbeat communication, and remote and (if supported) local disk logging. Separating these two planes means that resources used for traffic and security processing are not compromised by management activities.

FortiGate-6000 schematic



In the data plane, two DP3 load balancers use session-aware load balancing to distribute sessions from the front panel interfaces (port1 to 28) to Fortinet Processor Cards (FPCs). The DP3 processors communicate with the FPCs across the 3.2Tbps integrated switch fabric. Each FPC processes sessions load balanced to it. The FPCs send outgoing sessions back to the integrated switch fabric and then out the network interfaces to their destinations.

The NP6 processor in each FPC enhances network performance with fastpath acceleration that offloads communication sessions from the FPC CPU. The NP6 processor can also handle some CPU intensive tasks, like IPsec VPN encryption/decryption.

The CP9 processors in each FPC accelerate many common resource intensive security related processes such as SSL VPN, Antivirus, Application Control, and IPS.

The management plane includes the management board, base backplane, management interfaces, and HA heartbeat interfaces. Configuration and session synchronization between FPCs in a FortiGate-6000F occurs over the base backplane. In an HA configuration, configuration and session synchronization between the FortiGate-6000s in the cluster takes place over the HA1 and HA2 interfaces. Administrator logins, SNMP monitoring, remote logging to one or more FortiAnalyzers or syslog servers, and other management functions use the MGMT1, MGMT2, and MGMT3 interfaces. You can use the 10Gbps MGMT3 interface for additional bandwidth that might be useful for high bandwidth activities such as remote logging.

Interface groups and changing data interface speeds

Depending on the networks that you want to connect your FortiGate-6000 to, you may have to manually change the data interface speeds. The port1 to port20 data interfaces are divided into the following groups:

- port1 - port4
- port5 - port8
- port9 - port12
- port13 - port16
- port17 - port20
- port21 - port24

All of the interfaces in a group operate at the same speed. Changing the speed of an interface changes the speeds of all of the interfaces in the same group. For example, if you change the speed of port18 from 10Gbps to 25Gbps the speeds of port17 to port20 are also changed to 25Gbps.

The port25 to port28 interfaces are not part of an interface group. You can set the speed of each of these interfaces independently of the other three.

Another example, the default speed of the port1 to port24 interfaces is 10Gbps. If you want to install 25GigE transceivers in port1 to port24 to convert these data interfaces to connect to 25Gbps networks, you must enter the following from the CLI:

```
config system interface
  edit port1
    set speed 25000full
  next
  edit port5
    set speed 25000full
  next
  edit port9
    set speed 25000full
  next
  edit port13
    set speed 25000full
  next
  edit port17
    set speed 25000full
  next
  edit port21
    set speed 25000full
end
```

Every time you change a data interface speed, when you enter the `end` command, the CLI confirms the range of interfaces affected by the change. For example, if you change the speed of port5 the following message appears:

```
config system interface
  edit port5
    set speed 25000full
  end
port5-port8 speed will be changed to 25000full due to hardware limit.
Do you want to continue? (y/n)
```

FortiGate 6000F series hardware generations

Two generations of FortiGate 6000F series hardware are now available. Both generations support the same software features. Generation 2 has two hardware improvements:

- The FPCs include more memory.
- When connected to high-line AC power, generation 2 FortiGate 6000F series models provide 1+1 PSU redundancy. When connected to high-line AC power, each PSU provides 2000W, which is enough power to run the entire system including all FPCs.

For more information on FortiGate 6000F series generation 1 and generation 2, including supported firmware versions and how to determine the generation of your FortiGate 6000F series hardware, see the Fortinet Community article:

[Technical Tip: Information on FortiGate 6000F series Gen1 and Gen2.](#)

For more information on generation 1 and generation 2 AC PSUs, see [FortiGate-6000F AC power supply units \(PSUs\)](#).

The FortiGate-6001F is a generation 2 appliance.

FortiGate 6001F model licensing

The FortiGate 6001F model includes a total of ten FPCs, by default three of them are active. To increase throughput you can purchase subscription or perpetual licenses for each of the additional seven FPCs for a total of ten.

- A subscription license provides access to one or more FPCs for the term of the subscription. Subscription licenses are managed by FortiGuard.
- A perpetual license provides permanent access to one or more FPCs. Perpetual licenses are managed by applying a license key.

Subscription and perpetual licenses can be purchased from Fortinet. The package you get when you purchase the license includes instructions for applying the license to your FortiGate from the [Fortinet Support](#) website.

Applying your FortiGate 6001F FPC subscription license

To apply your FortiGate 6001F FPC subscription license, you must register your FortiGate 6001F and purchase a FPC subscription license for it. From the [Fortinet Support](#) website you can apply the FPC perpetual license to the FortiGate 6001F.

The subscription license will then be applied to the registered FortiGate 6001F through FortiGuard. It can take some time for the subscription license to be applied. From the FortiGate 6001F, you can enter the command `execute update-now` to start a FortiGuard update instead of waiting for the next scheduled update.

Activating your FortiGate 6001F FPC perpetual license

To activate your FortiGate 6001F FPC perpetual license, you must register your FortiGate 6001F and purchase a FPC perpetual license for it. From the [Fortinet Support](#) website you can apply the FPC perpetual license to the FortiGate 6001F and obtain your FPC perpetual license activation code or license key.

You can use the following command to apply your FPC perpetual license activation code or license key to activate FPCs for your FortiGate 6001F:

```
execute upd-perpetual-license <license-key>
```

After you enter this command, the FortiGate 6001F restarts and the licensed FPCs will be available.

Verifying your FortiGate 6001F FPC subscription and perpetual licenses

You can use the `diagnose load-balance status` command or the Cluster Status dashboard widget to view the number of FPCs that are available. These methods don't include any information about subscription or perpetual licenses, but they will show you how many FPCs are available, including the licensed FPCs.

You can enter the `get system status` command to see the number of FPCs licensed by your perpetual license and by your subscription license. In the following example, the FortiGate 6001F has a perpetual license for two FPCs and a subscription license for three FPCs.

```
get system status
.
.
.
FPC perpetual/subsription license: 2/3
```

You can also enter the following command to check the number of FPCs that you have subscription licenses for and the license expiration date. The following output shows a subscription license for three FPCs and that the license expires on June 26, 2025.

```
diagnose load-balance fpc-license
=====
MBD SN: <sn>
subscription-license : 3
subscription-license: Thu Jun 26 2025
```

What to expect when your subscription license expires

When the subscription license expires, the licensed FPCs are powered down and the output of the `diagnose load-balance fpc-license` command changes to the following:

```
diagnose load-balance fpc-license
=====
MBD SN: <sn>
subscription-license : 0
subscription-expiration: Expired
```

There is a one hour grace period between the time the license expires and when the FPCs are shut down. If you run the `diagnose load-balance fpc-license` command during that hour grace period, the FPCs shut down immediately.

Getting started with FortiGate 6000F series

This section is a quick start guide to connecting and configuring a FortiGate 6000F for your network.

Before using this chapter, your FortiGate 6000F should be mounted and connected to your grounding and power system. In addition, your FortiGate 6000Fs should be powered up and the front and back panel LEDs should indicate normal operation.

When your FortiGate 6000F is operating normally, the front panel LEDs should appear as follows.

LED	State
Status	Green
Alarm	Off
HA	Off
Power	Green
Connected network interfaces	Solid or flashing green.

During normal operation, the back panel PSU and fan try LEDs should all be solid green. This indicates that each component has power and is operating normally.

Once the system has initialized, you have a few options for connecting to the FortiGate 6000F GUI or CLI:

- Log in to the management board GUI by connecting MGMT1 or MGMT2 to your network and browsing to <https://192.168.1.99> or <https://192.168.2.99>.
- Log in to the management board CLI by connecting MGMT1 or MGMT2 to your network and using an SSH client to connect to 192.168.1.99 or 192.168.2.99.
- Log in to the management board CLI by connecting to the RJ-45 RS-232 CONSOLE port with settings: BPS: 9600, data bits: 8, parity: none, stop bits: 1, flow control: none.

The FortiGate 6000F ships with the following factory default configuration.

Option	Default Configuration
Administrator Account User Name	admin
Password	(none) For security reasons you should add a password to the admin account before connecting the FortiGate 6000F to your network. From the GUI, access the Global GUI and go to System > Administrators , edit the admin account, and select Change Password . From the CLI: <pre>config global config system admin edit admin set password <new-password> end</pre>

Option	Default Configuration
MGMT1 IP/Netmask	192.168.1.99/24
MGMT2 IP/Netmask	192.168.2.99/24

All configuration changes must be made from the management board GUI or CLI and not from individual FPCs.

All other management communication (for example, SNMP queries, remote logging, and so on) use the MGMT1 or MGMT2 interfaces and are handled by the management board.

Confirming startup status

Before verifying normal operation and making configuration changes and so on you should wait until the FortiGate 6000F is completely started up and synchronized. This can take a few minutes.

To confirm that the FortiGate 6000F is synchronized, you can refer to the Cluster Status dashboard widget. If the system is synchronized, the management board (slot 0) all of the FPCs (slots 1 to 6 or 1 to 10) should be visible and their **Configuration Status** should be **In Sync**. The Cluster Status dashboard widget also indicates if any FPCs are not synchronized.

Serial	Slot ID	Configuration Status	Role	Type	Sessions	Memory	CPU
F6KF31T018900143	0	In Sync	Primary	Management	62	10%	0%
FPC6KFT018901327	1	In Sync	Secondary	Dataplane	19	17%	0%
FPC6KFT018901372	2	In Sync	Secondary	Dataplane	17	17%	0%
FPC6KFT018901346	3	In Sync	Secondary	Dataplane	17	17%	0%
FPC6KFT018901574	4	In Sync	Secondary	Dataplane	17	17%	0%
FPC6KFT018901345	5	In Sync	Secondary	Dataplane	19	17%	0%



The FortiGate 6000F uses the Fortinet Security Fabric for communication and synchronization between the management board and the FPCs and for normal GUI operation. By default, the Security Fabric is enabled and must remain enabled for normal operation.

You can also view the **Sensor Information** dashboard widget to confirm that the system temperatures are normal and that all power supplies and fans are operating normally.

From the CLI you can use the `diagnose sys confsync status | grep in_sy` command to view the synchronization status of the management board and FPCs. If all of the FPCs are synchronized, each output line should include `in_sync=1`. If a line ends with `in_sync=0`, that FPC is not synchronized. The following example just shows a few output lines:

```
diagnose sys confsync status | grep in_sy
FPC6KF3E17900200, Secondary, uptime=5385.45, priority=119, slot_id=2:1, idx=2, flag=0x4, in_sync=1
F6KF313E17900031, Secondary, uptime=5484.74, priority=2, slot_id=1:0, idx=0, flag=0x10, in_sync=1
F6KF313E17900032, Primary, uptime=5488.57, priority=1, slot_id=2:0, idx=1, flag=0x10, in_sync=1
FPC6KF3E17900201, Secondary, uptime=5388.78, priority=120, slot_id=2:2, idx=2, flag=0x4, in_sync=1
F6KF313E17900031, Secondary, uptime=5484.74, priority=2, slot_id=1:0, idx=0, flag=0x10, in_sync=1
...
```

FortiGate 6000F and the Security Fabric

The FortiGate 6000F supports the Fortinet Security Fabric and all Security Fabric related features. You can set up the FortiGate 6000F to serve as the Security Fabric root and you can configure the FortiGate 6000F to join an existing Security Fabric. For more information see [Fortinet Security Fabric](#).

The FortiGate 6000F uses the Fortinet Security Fabric for communication and synchronization between the management board and the FPCs and for normal GUI operation. By default, the Security Fabric is enabled and must remain enabled for normal operation.

When adding a FortiGate 6000F to an existing security fabric, for normal operation you must authorize the FortiGate 6000F and all of the FPCs on the root FortiGate. Otherwise, the FortiGate 6000F management board will not be able to communicate with the FPCs.

You must also manually add a FortiAnalyzer to the FortiGate 6000F configuration, because the default FortiGate 6000F Security Fabric configuration has `configuration-sync` set to `local`, so the FortiGate 6000F doesn't get security fabric configuration settings, such as the FortiAnalyzer configuration, from the root FortiGate.

If the FortiGate 6000F is not joining a Security Fabric, Fortinet recommends that you do not change the Security Fabric configuration. You can verify the default Security Fabric configuration from the CLI:

```
config system csf
  set status enable
  set uid ''
  set upstream ''
  set source-ip 0.0.0.0
  set upstream-interface-select-method auto
  set upstream-port 8013
  set group-name "SLBC"
  set group-password <password>
  set accept-auth-by-cert enable
  set log-unification disable
  set authorization-request-type serial
  set fabric-workers 2
  set downstream-access disable
  set configuration-sync local
  set fabric-object-unification default
  set trusted-list ''
  set forticloud-account-enforcement enable
  set file-mgmt enable
  set file-quota 0
  set file-quota-warning 90
end
```

Configuration synchronization

When you log into the FortiGate 6000F GUI or CLI by connecting to the IP address of the MGMT1 or MGMT2 interface, or through a console connection, you are logging into the FortiGate 6000F management board (part of the FortiGate 6000F management plane). The management board is the FortiGate 6000F config-sync primary. All configuration changes must be made from the management board GUI or CLI. In an HA cluster, all configuration changes must be

made from the management board of the primary FortiGate 6000F. The management board synchronizes configuration changes to the FPCs and makes sure FPC configurations remain synchronized with the management board.

Once you have logged into the management board GUI or CLI and verified that the system is operating normally, you can view and change the configuration of your FortiGate 6000F just like any FortiGate.

For the FortiGate 6000F to operate normally, the configurations of the management board and all of the FPCs must be synchronized. You can use the information in the following sections to make sure that these configurations are synchronized.

Confirming that FortiGate-6000 components are synchronized

In addition to viewing configuration synchronization status from the Configuration Sync Monitor, you can use the following command to confirm that the configurations of the management board and the FPCs are synchronized:

```
diagnose sys confsync status
```

The command shows the HA and configuration synchronization (confsync) status of the management board and all of the FPCs. For the management board and each FPC, `in_sync=1` means the component is synchronized and can operate normally. If any component is out of sync, the command output will include `in_sync=0`. All components must be synchronized for the FortiGate-6000 to operate normally.



To confirm the configuration synchronization status of an HA cluster, see [Confirming that the FortiGate-6000 HA cluster is synchronized on page 61](#).

FPC confsync status

The `diagnose sys confsync status` command output begins with the confsync status for each FPC. In the following example for a FortiGate-6301F, the output begins with the confsync status if the FPC in slot 1. The two lines that begin with serial numbers and end with `in_sync=1` indicate that the FPC (serial number FPC6KFT018903332) is synchronized with the management board (serial number F6KF31T019900078) and the management board is synchronized with the FPC.

```
diagnose sys confsync status
...
Slot: 1  Module SN: FPC6KFT018903332
ELBC: svcgrp_id=1, chassis=1, slot_id=1
ELBC HB devs:
    elbc-ctrl/1: active=1, hb_count=10655
ELBC mgmt devs:
    elbc-b-chassis: mgmtip_set=1

zone: self_idx:1, primary_idx:0, ha_primary_idx:255, members:2
FPC6KFT018903332, Secondary, uptime=10654.18, priority=19, slot_id=1:1, idx=1, flag=0x4, in_sync=1
F6KF31T019900078, Primary, uptime=10697.67, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
    elbc-b-chassis: state=3(connected), ip=169.254.2.15, last_hb_time=10880.47, hb_nr=51017
...
```

Management board confsync status

The `diagnose sys confsync status` command output ends with the confsync status of the management board, which shows the configuration status between the management board and each of FPCs:

```
...
MBD SN: F6KF31T019900078
ELBC: svcgrp_id=1, chassis=1, slot_id=0

zone: self_idx:0, primary_idx:0, ha_primary_idx:255, members:7
F6KF31T019900078, Primary, uptime=11355.27, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
FPC6KFT018903307, Secondary, uptime=11309.71, priority=21, slot_id=1:3, idx=1, flag=0x24, in_sync=1
  elbc-b-chassis: state=3(connected), ip=169.254.2.5, last_hb_time=11431.39, hb_nr=54184
FPC6KFT018903310, Secondary, uptime=11314.01, priority=23, slot_id=1:5, idx=2, flag=0x24, in_sync=1
  elbc-b-chassis: state=3(connected), ip=169.254.2.7, last_hb_time=11431.43, hb_nr=54168
FPC6KFT018903327, Secondary, uptime=11313.17, priority=24, slot_id=1:6, idx=3, flag=0x24, in_sync=1
  elbc-b-chassis: state=3(connected), ip=169.254.2.8, last_hb_time=11431.38, hb_nr=54175
FPC6KFT018903328, Secondary, uptime=11312.37, priority=22, slot_id=1:4, idx=4, flag=0x24, in_sync=1
  elbc-b-chassis: state=3(connected), ip=169.254.2.6, last_hb_time=11431.36, hb_nr=54196
FPC6KFT018903332, Secondary, uptime=11312.10, priority=19, slot_id=1:1, idx=5, flag=0x24, in_sync=1
  elbc-b-chassis: state=3(connected), ip=169.254.2.3, last_hb_time=11431.38, hb_nr=54196
FPC6KFT018903333, Secondary, uptime=11313.94, priority=20, slot_id=1:2, idx=6, flag=0x24, in_sync=1
  elbc-b-chassis: state=3(connected), ip=169.254.2.4, last_hb_time=11431.38, hb_nr=54195
```

Viewing more details about FortiGate-6000 synchronization

If the output of the `diagnose sys configsync status` command includes `in_sync=0` entries, you can use the `diagnose sys confsync showcsum` command to view more details about the configuration checksums and potentially identify parts of the configuration that are not synchronized.

The `diagnose sys configsync showcsum` command shows HA and confsync debugzone and checksum information for the management board and the FPCs, beginning with the FPC in slot 1 and ending with the management board.

The following example shows the FPC in slot 1.

```
diagnose sys confsync showcsum
=====
Slot: 1  Module SN: FPC6KFT018903332
ha debugzone
global: e3 62 5b 5e 0e 99 3e 28 e2 27 72 d7 d4 16 a5 42
root: 3e af ce 65 00 45 f1 b6 c1 ae 65 40 0a 97 63 fc
mgmt-vdom: 2c 6e 41 c7 d0 15 34 e5 f1 c3 d9 9b 6f a4 fd 47
all: a2 89 77 7b 7f ad 38 b4 f3 16 53 17 f2 8b 60 61

ha checksum
global: e3 62 5b 5e 0e 99 3e 28 e2 27 72 d7 d4 16 a5 42
root: 3e af ce 65 00 45 f1 b6 c1 ae 65 40 0a 97 63 fc
mgmt-vdom: 2c 6e 41 c7 d0 15 34 e5 f1 c3 d9 9b 6f a4 fd 47
all: a2 89 77 7b 7f ad 38 b4 f3 16 53 17 f2 8b 60 61

confsync debugzone
global: be 96 26 6f a7 5d d1 d9 3f 8d 5f 45 46 80 9b 9d
root: 95 43 03 15 0b ce 2e 4e 55 e9 ec 37 65 47 d0 41
mgmt-vdom: c4 fc 49 b6 f1 ff c2 6d 9c bf 1e 5b 7d 5e 69 29
all: b3 f2 1a 4d fa fb b6 06 15 9a 42 17 ae 7e a0 be
```

```
confsync checksum
global: be 96 26 6f a7 5d d1 d9 3f 8d 5f 45 46 80 9b 9d
root: 95 43 03 15 0b ce 2e 4e 55 e9 ec 37 65 47 d0 41
mgmt-vdom: c4 fc 49 b6 f1 ff c2 6d 9c bf 1e 5b 7d 5e 69 29
all: b3 f2 1a 4d fa fb b6 06 15 9a 42 17 ae 7e a0 be
```

The example output includes four sets of checksums: a checksum for the global configuration, a checksum for each VDOM (in this case there are two VDOMs: root and mgmt-vdom), and a checksum for the complete configuration (all). You can verify that this FPC is synchronized because both sets of HA checksums match and both sets of confsync checksums match. Also as expected, the HA and confsync checksums are different.

If the management board and all of the FPCs in a standalone FortiGate-6000 have the same set of checksums, the management board and the FPCs in that FortiGate-6000 are synchronized.

If a FPC or the management board is out of sync, you can use the output of the `diagnose sys confsync status` command to determine what part of the configuration is out of sync. You could then take action to attempt to correct the problem or contact Fortinet Technical Support at <https://support.fortinet.com> for assistance.

A corrective action could be to restart of the component with the synchronization error. You could also try using the following command to re-calculate the checksums in case the sync error is just temporary:

```
diagnose sys confsync csum-recalculate
```

Cluster Status dashboard widget

The management board GUI dashboard includes the Cluster Status dashboard widget (one of the System widgets). You can use this widget to view the configuration synchronization status of your FortiGate 6000F management board and individual FPCs. The Cluster Status dashboard widget also displays the current number of sessions, memory usage, and CPU usage for the management board and the FPCs. This display allows you to separate management plane resource usage (management board) from data plane resource usage (FPCs). The Cluster Status dashboard widget is only available from the management board GUI.



When adding a FortiGate 6000F to an existing security fabric, for normal operation you must authorize the FortiGate 6000F and all of the FPCs on the root FortiGate. Otherwise, the FortiGate 6000F management board will not be able to communicate with the FPCs. No entries will appear on the Configuration Sync Monitor until all of the FPCs have been authorized with the root FortiGate.

Synchronization information includes the configuration status, role (primary or secondary), type (management or data), current number of sessions, current memory usage, current CPU usage, time since the last heartbeat was received from the component, the total number of heartbeat packets received from the component, and up time. You can also customize the columns that appear. If an FPC has failed, it will be removed from the list. If an FPC is out of synchronization this will be reflected in the Configuration Status list.

Serial	Slot ID	Configuration Status	Role	Type	Sessions	Memory	CPU
F6KF31T018900143	0	In Sync	Primary	Management	62	10%	0%
FPC6KFT018901327	1	In Sync	Secondary	Dataplane	19	17%	0%
FPC6KFT018901372	2	In Sync	Secondary	Dataplane	17	17%	0%
FPC6KFT018901346	3	In Sync	Secondary	Dataplane	17	17%	0%
FPC6KFT018901574	4	In Sync	Secondary	Dataplane	17	17%	0%
FPC6KFT018901345	5	In Sync	Secondary	Dataplane	19	17%	0%

You can hover your mouse cursor over the management board or the FPCs to view more detailed information including the hostname, serial number, firmware version, management port, CPU usage, memory usage, and session count.

From the pop up you can also select **Login** to log into the FPC's GUI using its management IP address and special port number. You can also select **Configure** to change the FPC's host name and some management-related options.

FortiGate 6000F dashboard widgets

The FortiGate 6000F includes custom dashboard widgets that provide extra or custom information for FortiGate 6000F systems.

Cluster status

Displays SLBC synchronization information, see [Cluster Status dashboard widget on page 22](#).

Resource Usage

The resource usage widgets show CPU usage, disk usage, log rate, memory usage, session setup rate, and the current number of sessions for the management board.

Sensor Information

The Sensor Information dashboard widget displays FortiGate 6000F temperature, power supply (PSU), and fan speed information. You can click on any item on the widget to display data collected by individual sensors.

Multi VDOM mode

By default, when you first start up a FortiGate 6000F it is operating in Multi VDOM mode. The default Multi VDOM configuration includes the **root** VDOM and a management VDOM named **mgmt-vdom**. The mgmt1, mgmt2, mgmt3, ha1, and ha2 interfaces are in mgmt-vdom and all of the data interfaces are in the root VDOM.

You cannot delete or rename mgmt-vdom. You also cannot remove interfaces from it or add interfaces to it. You can however, configure other settings such as routing required for management communication, interface IP addresses, and

so on. You can also add VLANs to the interfaces in mgmt-vdom and create a LAG that includes the mgmt1 and mgmt2 interfaces.

You can use the root VDOM for data traffic and you can also add more VDOMs for data traffic as required, depending on your Multi VDOM license.

FortiGate-6000 7.6.2 incompatibilities and limitations

FortiGate 6000F for FortiOS 7.6.2 has the following limitations and incompatibilities with FortiOS features:



The FortiGate 6000F uses the Fortinet Security Fabric for communication and synchronization between the management board and the FPCs and for normal GUI operation. By default, the Security Fabric is enabled and must remain enabled for normal operation.

Remote console limitations

Some console input may not function as expected. For example, when remotely connecting to an FPC console using Telnet, when viewing the BIOS menu, pressing the H key to display BIOS menu help does not always work as expected.

Default management VDOM

By default the FortiGate-6000 configuration includes a management VDOM named mgmt-vdom. The ha1, ha2, mgmt1, mgmt2, and mgmt3 interfaces are in mgmt-vdom and all other interfaces are in the root VDOM. For the FortiGate-6000 system to operate normally, mgmt-vdom must always be the management VDOM. You also must not remove interfaces from this VDOM. You can change the IP addresses of the interfaces in mgmt-vdom, allow the required management services, and add routes as required for management traffic.

You have full control over the configurations of other FortiGate-6000 VDOMs.

Maximum number of LAGs and interfaces per LAG

FortiGate-6000 systems support up to 16 link aggregation groups (LAGs). This includes both normal link aggregation groups and redundant interfaces. A FortiGate-6000 LAG can include up to 20 interfaces.

High Availability

Only the HA1 and HA2 interfaces are used for the HA heartbeat communication. For information on how to set up HA heartbeat communication using the HA1 and HA2 interfaces, see [Connect the HA1 and HA2 interfaces for HA heartbeat communication on page 54](#).

The following FortiOS HA features are not supported or are supported differently by FortiGate-6000:

- Active-active HA is not supported.
- The range for the HA `group-id` is 0 to 255.

- Failover logic for FortiGate-6000 HA is the same as FGCP for other FortiGate clusters.
- HA heartbeat configuration is specific to FortiGate-6000 systems and differs from standard HA.
- FortiGate-6000 HA does not support the `route-wait` and `route-hold` options for tuning route synchronization between FortiGate-6000s.
- VLAN monitoring using the `config system ha-monitor` command is not supported.
- FortiGate-6000 HA does not support using the HA `session-sync-dev` option. Instead, session synchronization traffic uses the HA1 and HA2 interfaces, separating session sync traffic from data traffic.
- For information about virtual clustering limitations, see [Limitations of FortiGate-6000 virtual clustering on page 69](#) and [Virtual clustering VLAN/VDOM limitation on page 69](#).

FortiOS features that are not supported by FortiGate-6000 7.6.2

The following mainstream FortiOS features are not supported by the FortiGate-6000:

- Hardware switch.
- DLP archiving.
- GRE tunneling is only supported after creating a load balance flow rule, for example:

```
config load-balance flow-rule
  edit 0
    set status enable
    set vlan 0
    set ether-type ip
    set protocol gre
    set action forward
    set forward-slot master
    set priority 3
  end
```
- Only the FortiGate 6001F, FortiGate 6301F, and FortiGate 6501F support hard disk features such as disk logging, and GUI-based packet sniffing.
- The FortiGate 6000F platform, including the FortiGate 6001F, FortiGate 6301F, and FortiGate 6501F do not support quarantining files to the internal hard disks. Instead you must set the quarantine function to quarantine files to FortiAnalyzer.
- The management interfaces (mgmt1-3) do not support device detection for the networks they are connected to.
- The FortiGate-6000 does not support configuring dedicated management interfaces using the `config system dedicated-mgmt` command or by enabling the `dedicated-to management interface` option. The purpose of the dedicated management interface feature is to add a routing table just for management connections. This functionality is supported by the FortiGate-6000 management VDOM (mgmt-vdom) that has its own routing table and contains all of the FortiGate-6000 management interfaces.
- The FortiOS `session-ttl` option `never` (which means no session timeout) is only supported if the `dp-load-distribution-method` is set to `src-dst-ip-sport-dport` (the default) or `src-dst-ip` and the firewall policy that accepts the session does not perform NAT. If any other load distribution method is used, or if NAT is enabled, the DP session timer will terminate the session according to the DP processor session timer. For more information about the `never` option, see [No session timeout](#).
- Enabling the system settings option `tcp-session-without-syn` and configuring a firewall policy to accept sessions without syn packets allows FortiOS to add entries to its session table for sessions that do not include SYN packets. These sessions can only be load balanced by the DP processor if the `dp-load-distribution-method` is set to `src-dst-ip-sport-dport` (default) or `src-dst-ip`. If any other load distribution method is used, the sessions will be dropped. As well, the DP processor cannot load balance these sessions if they are accepted by a firewall policy with NAT enabled.

- The `source-ip` option for management services (for example, logging, SNMP, connecting to FortiSandbox) that use interfaces in the `mgmt-vdom` is not supported and has been removed from the CLI.
- The `config vpn ssl settings option tunnel-addr-assigned-method` is now available again. This option had been removed from the CLI in a previous release because setting this option to `first-available` and configuring multiple IP pools can reduce FortiGate 6000F SSL VPN load balancing performance. However, some users may want the ability to use multiple IP pools for their SSL VPN configuration, even if performance is reduced. So the change has been reverted.
- The `encryption` option of the `config system standalone-cluster` command does not encrypt session synchronization traffic. Enabling this option has no effect.

IPsec VPN tunnels terminated by the FortiGate 6000F

For a list of IPsec VPN features not supported by FortiGate 6000F, see [IPsec VPN load balancing on page 47](#).

Traffic shaping and DDoS policies

Each FPC applies traffic shaping and DDoS quotas independently. Because of load-balancing, this may allow more traffic than expected.

FortiGuard web filtering and spam filtering queries

The FortiGate-6000 sends all FortiGuard web filtering and spam filtering rating queries through a management interface from the `mgmt-vdom` VDOM.

Web filtering quotas

On a VDOM operating with the **Inspection Mode** set to **Proxy**, you can go to **Security Profiles > Web Filter** and set up **Category Usage Quotas**. Each FPC has its own quota, and the FortiGate-6000 applies quotas per FPC and not per the entire FortiGate-6000 system. This could result in quotas being exceeded if sessions for the same user are processed by different FPCs.

Special notice for new deployment connectivity testing

Only the management board can successfully ping external IP addresses. During a new deployment, while performing connectivity testing from the Fortigate-6000, make sure to run `execute ping` tests from the management board and not from an FPC. See [Using data interfaces for management traffic on page 98](#) for information about changes to this limitation.

Display the process name associated with a process ID

You can use the following command to display the process name associated with a process ID (PID):

```
diagnose sys process nameof <pid>
```

Where `<pid>` is the process ID.

Managing individual FortiGate 6000F management boards and FPCs

You can manage individual FPCs using special management port numbers, FPC consoles, or the `execute load-balance slot manage` command. You can also use the `execute ha manage` command to log in to the other FortiGate 6000F in an HA configuration.

Special management port numbers

You may want to connect to individual FPCs to view status information or perform a maintenance task, such as installing firmware or performing a restart. You can connect to the GUI or CLI of individual FPCs (or the management board) using the MGMT1 interface IP address with a special port number.



You can use the `config load-balance setting slbc-mgmt-intf` command to change the management interface used. The default is `mgmt1` and it can be changed to `mgmt2`, or `mgmt3`.

To enable using the special management port numbers to connect to individual FPCs, set `slbc-mgmt-intf` to an interface that is connected to a network, has a valid IP address, and has management or administrative access enabled. To block access to the special management port numbers you can set `slbc-mgmt-intf` to an interface that is not connected to a network, does not have a valid IP address, or has management or administrative access disabled.

For example, if the MGMT1 interface IP address is 192.168.1.99 you can connect to the GUI of the first FPC (the FPC in slot 1) by browsing to :

```
https://192.168.1.99:44301
```

The special port number (in this case, 44301) is a combination of the service port (for HTTPS, the service port is 443) and the FPC slot number (in this example, 01).

You can view the special HTTPS management port number for and log in to the GUI of an FPC from the Configuration Sync Monitor.

The following table lists the special ports you can use to connect to individual FPCs or the management board using common management protocols. The FortiGate 6300F and 6301F have 7 slots (0 to 6) and the FortiGate 6500F and 6501F have 11 slots (0 to 10). Slot 0 is the management board (MBD) slot. Slots 1 to 10 are FPC slots.



You can't change the special management port numbers. Changing configurable management port numbers, for example the HTTPS management port number (which you might change to support SSL VPN), does not affect the special management port numbers.

FortiGate 6000F special management port numbers

Slot Address	HTTP (80)	HTTPS (443)	Telnet (23)	SSH (22)	SNMP (161)
Slot 0, (MBD)	8000	44300	2300	2200	16100
Slot 1 (FPC01)	8001	44301	2301	2201	16101
Slot 2 (FPC02)	8002	44302	2302	2202	16102
Slot 3 (FPC03)	8003	44303	2303	2203	16103
Slot 4 (FPC04)	8004	44304	2304	2204	16104
Slot 5 (FPC05)	8005	44305	2305	2205	16105
Slot 6 (FPC06)	8006	44306	2306	2206	16106
Slot 7 (FPC07)	8007	44307	2307	2207	16107
Slot 8 (FPC08)	8008	44308	2308	2208	16108
Slot 9 (FPC09)	8009	44309	2309	2209	16109
Slot 10 (FPC10)	8010	44310	2310	2210	16110

For example, to connect to the CLI of the FPC in slot 3 using SSH, you would connect to `ssh://192.168.1.99:2203`.

To verify which slot you have logged into, the GUI header banner and the CLI prompt shows the current hostname. The System Information dashboard widget also shows the host name and serial number. The CLI prompt also shows slot address in the format `<hostname> [<slot address>] #`.

Logging in to different FPCs allows you to use the FortiView or Monitor GUI pages to view the activity on that FPC. You can also restart the FPC from its GUI or CLI. Even though you can log in to different FPCs, you can only make configuration changes from the management board.

HA mode special management port numbers

In an HA configuration consisting of two FortiGate 6000Fs in an HA cluster, you can connect to individual FPCs or to the management board in chassis 1 (chassis ID = 1) using the same special port numbers as for a standalone FortiGate 6000F.

You use different special port numbers to connect to individual FPCs or the management board in the FortiGate 6000F with chassis ID 2 (chassis ID = 2).

FortiGate 6000F special management port numbers (chassis ID = 2)

Slot Address	HTTP (80)	HTTPS (443)	Telnet (23)	SSH (22)	SNMP (161)
Slot 0, (MBD)	8020	44320	2320	2220	16120
Slot 1 (FPC01)	8021	44321	2321	2221	16121
Slot 2 (FPC02)	8022	44322	2322	2222	16122

Slot Address	HTTP (80)	HTTPS (443)	Telnet (23)	SSH (22)	SNMP (161)
Slot 3 (FPC03)	8023	44323	2323	2223	16123
Slot 4 (FPC04)	8024	44324	2324	2224	16124
Slot 5 (FPC05)	8025	44325	2325	2225	16125
Slot 6 (FPC06)	8026	44326	2326	2226	16126
Slot 7 (FPC07)	8027	44327	2327	2227	16127
Slot 8 (FPC08)	8028	44328	2328	2228	16128
Slot 9 (FPC09)	8029	44329	2329	2229	16129
Slot 10 (FPC10)	8030	44330	2330	2230	16130

Connecting to individual FPC consoles

From the management board CLI, you can use the `execute system console-server` command to access individual FPC consoles. Console access can be useful for troubleshooting. For example, if an FPC does not boot properly, you can use console access to view the state of the FPC and enter commands to fix the problem or restart the FPC.

From the console, you can also perform BIOS-related operations, such as rebooting the FPC, interrupting the boot process, and installing new firmware.

For example, from the management board CLI, use the following command to log in to the console of the FPC in slot 3:

```
execute system console-server connect 3
```

Authenticate to log in to the console and use CLI commands to view information, make changes, or restart the FPC. When you are done, use **Ctrl-X** to exit from the console back to the management board CLI. Using **Ctrl-X** may not work if you are accessing the CLI console from the GUI. Instead you may need to log out of the GUI and then log in again.

Also, from the management board CLI you can use the `execute system console-server showline` command to list any active console server sessions. Only one console session can be active for each FPC, so before you connect to an FPC console, you can use the following command to verify whether or not there is an active console session. The following command output shows an active console session with the FPC in slot 4:

```
execute system console-server showline
MB console line connected - 1
Telnet-to-console line connected - 4
```

To clear an active console session, use the `execute system console-server clearline` command. For example, to clear an active console session with the FPC in slot 4, enter:

```
execute system console-server clearline 4
```



In an HA configuration, the `execute system console-server` commands only allow access to FPCs in the FortiGate 6000F that you are logged into. You can't use this command to access FPCs in the other FortiGate 6000F in an HA cluster

Connecting to individual FPC CLIs

From the management board CLI you can use the following command to log into the CLI of individual FPCs:

```
execute load-balance slot manage {<chassis>.slot | slot}
```

Where:

<slot> is the slot number of the component that you want to log in to. The management board is in slot 0 and the FPC slot numbers start at 1. In an FGCP HA cluster you must also include the <chassis> number which can be 1 or 2.

When connected to the CLI of a FPC, you can view information about the status or configuration of the FPC, restart the FPC, or perform other operations. You should not change the configuration of individual FPCs because this can cause configuration synchronization errors.

Use the `exit` command to end the session and return to the CLI from which you ran the original command.

Connecting to individual FPC CLIs of the secondary FortiGate-6000 in an HA configuration

From the management board of the primary FortiGate-6000, you can use the following command to log in to the management board of the secondary FortiGate-6000:

```
execute ha manage <id>
```

Where <id> is the ID of the other FortiGate-6000 in the cluster. From the primary FortiGate-6000, use an ID of 0 to log into the secondary FortiGate-6000. From the secondary FortiGate-6000, use an ID of 1 to log into the primary FortiGate-6000. You can enter the ? to see the list of IDs that you can connect to.

After you have logged in, you can manage the secondary FortiGate-6000 from the management board or you can use the `execute load-balance slot manage` command to connect to the CLIs of different FPCs in the secondary FortiGate-6000.

Performing other operations on individual FPCs

You can use the following commands to restart, power off, power on, or perform an NMI reset on individual FPCs while logged into the management board CLI:

```
execute load-balance slot {nmi-reset | power-off | power on | reboot | set-primary-worker} <slots>
```

Where <slots> can be one or more slot numbers or slot number ranges separated by commas. Do not include spaces.

For example, to shut down the FPCs in slots 2, and 4 to 6 enter:

```
execute load-balance slot power-off 2,4-6
```

Load balancing and flow rules

This chapter provides an overview of how FortiGate-6000 Session-Aware Load Balancing (SLBC) works and then breaks down the details and explains why you might want to change some load balancing settings.

FortiGate-6000 SLBC works as follows.

1. The FortiGate-6000 directs all traffic that does not match a load balancing flow rule to the DP3 processors.
If a session matches a flow rule, the session skips the DP3 processors and is directed according to the action setting of the flow rule. Default flow rules send traffic that can't be load balanced to the primary FPC. See [Default flow rules for traffic that cannot be load balanced on page 33](#).

2. The DP3 processors load balance TCP, UDP, SCTP, and (if enabled) ESP (IPsec) sessions among the FPCs according to the load balancing method set by the `dp-load-distribution-method` option of the `config load-balance setting` command.

The DP3 processors load balance ESP (IPsec) sessions that use static routes if IPsec VPN load balancing is enabled. If IPsec VPN load balancing is disabled, the DP3 processors send ESP (IPsec) sessions to the primary FPC. For more information about IPsec VPN load balancing, see [Configuring IPsec VPN load balancing on page 47](#).

The DP3 processors load balance ICMP sessions according to the load balancing method set by the `dp-icmp-distribution-method` option of the `config load-balance setting` command. See [ICMP load balancing on page 37](#).

The DP3 processors load balance GTP-U sessions if GTP load balancing is enabled. If GTP load balancing is disabled, the DP3 processors send GTP sessions to the primary FPC. For more information about GTP load balancing, see [Enabling GTP load balancing on page 34](#).

The DP3 processors load balance PFCP-controlled GTP-U sessions if PFCP load balancing is enabled. If PFCP load balancing is disabled, the DP3 processors send PFCP-controlled GTP-U sessions to the primary FPC. For more information about PFCP load balancing, see [PFCP load balancing on page 36](#).

To support ECMP, you can change how the DP3 processors manage session tables, see [ECMP support on page 90](#).

3. The DP3 processors send other sessions that cannot be load balanced to the primary FPC.

Setting the load balancing method

The FortiGate-6000 DP3 processor load balances or distributes sessions based on the load balancing method set by the following command:

```
config load-balance setting
  set dp-load-distribution-method {to-master | src-ip | dst-ip | src-dst-ip | src-ip-sport
    | dst-ip-dport | src-dst-ip-sport-dport}
end
```

The default load balancing method, `src-dst-ip-sport-dport`, distributes sessions across all FPCs according to their source and destination IP address, source port, and destination port. This load balancing method represents true session-aware load balancing. Session aware load balancing takes all session information into account when deciding where to send new sessions and where to send additional packets that are part of an already established session.

For information about the other load balancing methods, see [config load-balance setting on page 119](#).

Determining the primary FPC

You can determine which FPC is operating as the primary FPC by using the `diagnose load-balance status` command.

The following example `diagnose load-balance status` output for a FortiGate-6301F shows that the FPC in slot 1 is the primary FPC. The command output also shows the status of all of the FPCs in the FortiGate-6301F.

```
diagnose load-balance status
=====
MBD SN: F6KF313E17900032
  Primary FPC Blade: slot-1

Slot 1: FPC6KF3E1790020
  Status:Working  Function:Active
  Link:          Base: Up          Fabric: Up
  Heartbeat: Management: Good    Data: Good
  Status Message:"Running"
Slot 2: FPC6KF3E17900201
  Status:Working  Function:Active
  Link:          Base: Up          Fabric: Up
  Heartbeat: Management: Good    Data: Good
  Status Message:"Running"
Slot 3: FPC6KF3E17900207
  Status:Working  Function:Active
  Link:          Base: Up          Fabric: Up
  Heartbeat: Management: Good    Data: Good
  Status Message:"Running"
Slot 4: FPC6KF3E17900219
  Status:Working  Function:Active
  Link:          Base: Up          Fabric: Up
  Heartbeat: Management: Good    Data: Good
  Status Message:"Running"
Slot 5: FPC6KF3E17900235
  Status:Working  Function:Active
  Link:          Base: Up          Fabric: Up
  Heartbeat: Management: Good    Data: Good
  Status Message:"Running"
Slot 6: FPC6KF3E17900169
  Status:Working  Function:Active
  Link:          Base: Up          Fabric: Up
  Heartbeat: Management: Good    Data: Good
  Status Message:"Running"
```

Flow rules for sessions that cannot be load balanced

Some traffic types cannot be load balanced. Sessions for traffic types that cannot be load balanced should normally be sent to the primary FPC by configuring flow rules for that traffic. You can also configure flow rules to send traffic that

cannot be load balanced to specific FPCs.

Create flow rules using the `config load-balance flow-rule` command. The default configuration uses this command to send a number of traffic types that cannot be load balanced to the primary FPC. You can view the default configuration of the `config load-balance flow-rule` command to see how this is all configured, or see [Default flow rules for traffic that cannot be load balanced on page 33](#).

For example, the following configuration sends BGP source and destination sessions to the primary FPC:

```
config load-balance flow-rule
  edit 3
    set status enable
    set vlan 0
    set ether-type ip
    set protocol tcp
    set src-l4port 179-179
    set dst-l4port 0-0
    set tcp-flag any
    set action forward
    set forward-slot master
    set priority 5
    set comment "bgp src"
  next
  edit 4
    set status enable
    set vlan 0
    set ether-type ip
    set protocol tcp
    set src-l4port 0-0
    set dst-l4port 179-179
    set tcp-flag any
    set action forward
    set forward-slot master
    set priority 5
    set comment "bgp dst"
end
```

Default flow rules for traffic that cannot be load balanced

The default `configure load-balance flow-rule` command contains the recommended default flow rules that control how the FortiGate 6000F handles traffic types that cannot be load balanced. All of the default flow rules identify the traffic type using the options available in the command and direct matching traffic to the primary FPC (`action` set to `forward` and `forward-slot` set to `master`). The default flow rules also include a comment that identifies the traffic type. Most of the flow rules in the default configuration are enabled and are intended to send common traffic types that cannot be load balanced to the primary FPC.

The default configuration also includes disabled flow rules for Kerberos and PPTP traffic. Normally, you would only need to enable these flow rules if you know that your FortiGate 6000F will be handling these types of traffic.



You can find text files containing the FortiGate 6000F default flow rules in the same folders as firmware images on the [Fortinet Support Download Firmware Images page](#).

GTP load balancing

You can use the information in this section to optimize FortiGate 6000F GTP performance.

Enabling GTP load balancing

You can use the following load balancing command to enable or disable FortiGate 6000F GTP-U load balancing.

```
config load-balance setting
  config gtp-load-balance {disable | enable}
end
```

The following flow rule is also available to direct GTP-C traffic to the primary FPC.

```
config load-balance flow-rule
  edit 17
  set status disable
  set ether-type ipv4
  set src-addr-ipv4 0.0.0.0 0.0.0.0
  set dst-addr-ipv4 0.0.0.0 0.0.0.0
  set protocol udp
  set src-l4port 0-0
  set dst-l4port 2123-2123
  set action forward
  set forward-slot master
  set priority 5
  set comment "gtp-c to primary blade"
  next
end
```

The recommended configuration for optimal GTP-C tunnel setup and GTP-U throughput performance is to enable `gtp-load-balance` and disable the GTP-C flow rule. In this configuration, both GTP-C and GTP-U traffic is load balanced among all of the FPCs:

```
config load-balance setting
  config gtp-load-balance enable
end
config load-balance flow-rule
  edit 17
  set status disable
  set ether-type ipv4
  set src-addr-ipv4 0.0.0.0 0.0.0.0
  set dst-addr-ipv4 0.0.0.0 0.0.0.0
  set protocol udp
  set src-l4port 0-0
  set dst-l4port 2123-2123
  set action forward
  set forward-slot master
  set priority 5
  set comment "gtp-c to primary blade"
  next
end
```

If you want GTP-C traffic to only be processed by the primary FPC, you can edit the GTP-C flow rule and set `status` to `enable`. When enabled, this flow rule sends all GTP-C traffic to the primary FPC. Enabling this flow rule can reduce GTP

performance, since all GTP-C tunnel setup sessions will be done by the primary FPC and not distributed among all of the FPCs.

```
config load-balance flow-rule
  edit 17
    set status enable
    set ether-type ipv4
    set src-addr-ipv4 0.0.0.0 0.0.0.0
    set dst-addr-ipv4 0.0.0.0 0.0.0.0
    set protocol udp
    set src-l4port 0-0
    set dst-l4port 2123-2123
    set action forward
    set forward-slot master
    set priority 5
    set comment "gtp-c to primary blade"
  next
end
```

GTP-U load balancing may not distribute sessions evenly among all of the FPCs. Its common in many 4G networks to have just a few SGWs. Similar configurations with very few servers may also be used in other GTP implementations. If the FortiGate 6000F receives GTP traffic from a very few servers, the GTP traffic will have very few source and destination IP addresses and TCP/IP ports. Since SLBC load balancing is based on source and destination IP addresses and TCP ports, its possible that sessions will not be distributed evenly among the FPCs. In fact, most GTP-U traffic could be processed by a limited number of FPCs.

Enabling GTP-U load balancing still distributes sessions and improves performance, but performance gains from enabling GTP-U load balancing may not be as high as anticipated.

GTP with dynamic source port allocation

GTP supports dynamically allocating UDP source ports. Dynamic source port allocation can help balance the load on the network.

To support GTP with dynamic source port allocation on a FortiGate 6000F, the recommended configuration is to enable the standard GTP-C flow rule to send all GTP-C traffic to the primary FPC. This flow rule sends all traffic with destination port 2123 to the primary FPC.

```
config load-balance flow-rule
  edit 17
    set status enable
    set ether-type ipv4
    set src-addr-ipv4 0.0.0.0 0.0.0.0
    set dst-addr-ipv4 0.0.0.0 0.0.0.0
    set protocol udp
    set src-l4port 0-0
    set dst-l4port 2123-2123
    set action forward
    set forward-slot master
    set priority 5
    set comment "gtp-c to primary blade"
  next
end
```

In addition, you must create a second flow rule to send all GTP-C traffic with source port 2123 to the primary FPC. For example:

```

config load-balance flow-rule
  edit 23
  set status enable
  set ether-type ipv4
  set src-addr-ipv4 0.0.0.0 0.0.0.0
  set dst-addr-ipv4 0.0.0.0 0.0.0.0
  set protocol udp
  set src-l4port 2123-2123
  set dst-l4port 0-0
  set action forward
  set forward-slot master
  set priority 5
  set comment "gtp-c src port to primary blade"
  next
end

```



You can also support GTP with dynamic source port allocation by setting the `dp-load-distribution-method` to `src-dst-ip`:

```

config load-balance setting
  set dp-load-distribution-method src-dst-ip
end

```

Enabling GTP load balancing using the following command is supported for GTP with dynamic source port allocation:

```

config load-balance setting
  set gtp-load-balance {disable | enable}
end

```

GTP load balancing and fabric channel usage

On a FortiGate 6000F, when GTP load balancing is enabled, GTP tunnels are synchronized over the fabric channel backplane (also called the data channel). The fabric channel is also used for SLBC session synchronization. On a busy FortiGate 6000F that is also load balancing GTP tunnels, the system may experience more lost SLBC heartbeats than normal.

To avoid missed heartbeats, you should increase the `max-miss-heartbeats` load balancing setting.

For example, when GTP load balancing is enabled, Fortinet recommends setting the `max-miss-heartbeats` to 40.

```

config load-balance setting
  set max-miss-heartbeats 40
  set gtp-load-balance enable
end

```

PFCP load balancing

FortiGate 6000F includes support for load balancing the Packet Forwarding Control Protocol (PFCP). PFCP is an addition to 3GPP that provides 4G Control plane and User Plane Separation (CUPS) and 5G signaling evolution. When PFCP is used as the control plane, the user plane employs GTP-U encapsulation. PFCP takes many of the roles that are provided by GTP-C in 3G/4G networks and provides session awareness and tracking of GTP-U user plane traffic while also providing control plane initiation.

FortiGate 6000F PFCP support includes supporting PFCP session synchronization for FGCP HA.

You can use the following command to enable or disable FortiGate 6000F PFCP load balancing.

```
config load-balance setting
  set pfcpl-load-balance {disable | enable}
end
```

The following flow rule is also available to direct PFCP control plane traffic to the primary FPC.

```
edit 21
  set status disable
  set vlan 0
  set ether-type ipv4
  set src-addr-ipv4 0.0.0.0 0.0.0.0
  set dst-addr-ipv4 0.0.0.0 0.0.0.0
  set protocol udp
  set src-l4port 0-0
  set dst-l4port 8805-8805
  set action forward
  set forward-slot master
  set priority 5
  set comment "pfcpl to primary blade"
end
```

By default, both of these configurations are disabled and PFCP control plane and user plane traffic is not load balanced. The DP sends all PFCP control plane and user plane traffic to the primary FPC.

To load balance PFCP user plane traffic to multiple FPCs, you can set `pfcpl-load-balance` to `enable`. This also enables the PFCP flow rule. PFCP user plane traffic is then load balanced across all FPCs while PFCP control plane traffic is still handled by the primary FPC. This is the recommended configuration for load balancing PFCP traffic.

ICMP load balancing

You can use the following option to configure load balancing for ICMP sessions:

```
config load-balance setting
  set dp-icmp-distribution-method {to-master | src-ip | dst-ip | src-dst-ip | derived}
end
```

The default setting is `to-master` and all ICMP traffic is sent to the primary FPC.

If you want to load balance ICMP sessions to multiple FPCs, you can select one of the other options. You can load balance ICMP sessions by source IP address, by destination IP address, or by source and destination IP address.

You can also select `derived` to load balance ICMP sessions using the `dp-load-distribution-method` setting. Since port-based ICMP load balancing is not possible, if `dp-load-distribution-method` is set to a load balancing method that includes ports, ICMP load balancing will use the equivalent load balancing method that does not include ports. For example, if `dp-load-distribution-method` is set to the `src-dst-ip-sport-dport` (the default) then ICMP load balancing will use `src-dst-ip` load balancing.

Load balancing TCP, UDP, and ICMP sessions with fragmented packets

If your FortiGate-6000 receives fragmented TCP, UDP, or ICMP packets, you can use the following configuration to make sure the Internal Switch Fabric (ISF) handles them correctly.

```
config load-balance setting
  set dp-fragment-session enable
  set sw-load-distribution-method src-dst-ip
end
```

With this configuration, when the DP3 processor receives a header fragment packet, if a matching session is found, the DP3 processor creates an additional fragment session matching the source-ip, destination-ip, and IP identifier (IPID) of the header fragment packet. Subsequent non-header fragments will match this fragment session and be forwarded to the same FPC as the header fragment.

If `dp-fragment-session` is disabled (the default), handling fragmented packets is less efficient because the DP3 processor broadcasts all non-header fragmented TCP, UDP, or ICMP packets to all FPCs. FPCs that also received the header fragments of these packets re-assemble the packets correctly. FPCs that did not receive the header fragments discard the non-header fragments.



If `sw-load-distribution-method` is set to `src-dst-ip-sport-dport`, fragmented packets may be dropped. Changing the load distribution method to `src-dst-ip` may lower performance because regular traffic may not be optimally load balanced. You can experiment with enabling and disabling `dp-fragment-session` and changing `sw-load-distribution-method` to determine the configuration that produces the best results for your network's traffic.

The age of the fragment session can be controlled using the following command:

```
config system global
  set dp-fragment-timer <timer>
end
```

The default `<timer>` value is 120 seconds. The range is 1 to 65535 seconds.

Adding flow rules to support DHCP relay

The FortiGate-6000 default flow rules may not handle DHCP relay traffic correctly.

The default configuration includes the following flow rules for DHCP traffic:

```
config load-balance flow-rule
  edit 7
    set status enable
    set vlan 0
    set ether-type ipv4
    set src-addr-ipv4 0.0.0.0 0.0.0.0
    set dst-addr-ipv4 0.0.0.0 0.0.0.0
    set protocol udp
    set src-l4port 67-67
    set dst-l4port 68-68
```

```

    set action forward
    set forward-slot master
    set priority 5
    set comment "dhcpv4 server to client"
next
edit 8
    set status enable
    set vlan 0
    set ether-type ipv4
    set src-addr-ipv4 0.0.0.0 0.0.0.0
    set dst-addr-ipv4 0.0.0.0 0.0.0.0
    set protocol udp
    set src-l4port 68-68
    set dst-l4port 67-67
    set action forward
    set forward-slot master
    set priority 5
    set comment "dhcpv4 client to server"
end

```

These flow rules handle traffic when the DHCP client sends requests to a DHCP server using port 68 and the DHCP server responds using port 67. However, if DHCP relay is involved, requests from the DHCP relay to the DHCP server and replies from the DHCP server to the DHCP relay both use port 67. If this DHCP relay traffic passes through the FortiGate-6000 you must add a flow rule similar to the following to support port 67 DHCP traffic in both directions (the following example uses `edit 0` to add the DHCP relay flow using the next available flow rule index number):

```

config load-balance flow-rule
    edit 0
        set status enable
        set vlan 0
        set ether-type ipv4
        set src-addr-ipv4 0.0.0.0 0.0.0.0
        set dst-addr-ipv4 0.0.0.0 0.0.0.0
        set protocol udp
        set src-l4port 67-67
        set dst-l4port 67-67
        set action forward
        set forward-slot master
        set priority 5
        set comment "dhcpv4 relay"
    next

```

The default configuration also includes the following flow rules for IPv6 DHCP traffic:

```

edit 13
    set status enable
    set vlan 0
    set ether-type ipv6
    set src-addr-ipv6 ::/0
    set dst-addr-ipv6 ::/0
    set protocol udp
    set src-l4port 547-547
    set dst-l4port 546-546
    set action forward
    set forward-slot master
    set priority 5
    set comment "dhcpv6 server to client"
next

```

```
edit 14
  set status enable
  set vlan 0
  set ether-type ipv6
  set src-addr-ipv6 ::/0
  set dst-addr-ipv6 ::/0
  set protocol udp
  set src-l4port 546-546
  set dst-l4port 547-547
  set action forward
  set forward-slot master
  set priority 5
  set comment "dhcpv6 client to server"
next
```

These flow rules handle traffic when the IPv6 DHCP client sends requests to a DHCP server using port 547 and the DHCP server responds using port 546. However, if DHCP relay is involved, requests from the DHCP relay to the DHCP server and replies from the DHCP server to the DHCP relay both use port 547. If this DHCP relay traffic passes through the FortiGate 6000F you must add a flow rule similar to the following to support port 547 DHCP traffic in both directions (the following example uses `edit 0` to add the DHCP relay flow using the next available flow rule index number):

```
config load-balance flow-rule
  edit 0
    set status enable
    set vlan 0
    set ether-type ipv6
    set src-addr-ipv6 ::/0
    set dst-addr-ipv6 ::/0
    set protocol udp
    set src-l4port 547-547
    set dst-l4port 547-547
    set action forward
    set forward-slot master
    set priority 5
    set comment "dhcpv6 relay"
  next
```

Flow rules to support multihop BFD (MBFD)

The FortiGate 6000F supports Multihop BFD for normal traffic and over IPsec VPN tunnels that are terminated by the FortiGate 6000F (see [BFD for multihop path for BGP](#)).

The multihop control protocol uses TCP and UDP traffic on port 4784. Multihop control traffic is not load balanced by DP processors. Instead, a flow rule is used to send all multihop control traffic to a single FPC.

The following flow rule has been added to the default flow rules for traffic that cannot be load balanced to send all multihop control traffic to the primary FPC. This flow rule should be enabled if you configure multihop BFD support on your FortiGate 6000F.

```
config load-balance flow-rule
  edit 22
    set status disable
    set vlan 0
    set ether-type ip
    set protocol udp
```

```
set src-l4port 0-0
set dst-l4port 4784-4784
set action forward
set forward-slot master
set priority 5
set comment "Flow Rule for Multihop BFD"
end
```

Flow rules to support IP multicast

IPv4 and IPv6 Multicast traffic cannot be load balanced by DP processors and instead is sent to the primary FPC. This is controlled by the following default flow rules:

```
config load-balance flow-rule
edit 15
set status enable
set vlan 0
set ether-type ipv4
set src-addr-ipv4 0.0.0.0 0.0.0.0
set dst-addr-ipv4 224.0.0.0 240.0.0.0
set protocol any
set action forward
set forward-slot master
set priority 5
set comment "ipv4 multicast"
next
edit 16
set status enable
set vlan 0
set ether-type ipv6
set src-addr-ipv6 ::/0
set dst-addr-ipv6 ff00::/8
set protocol any
set action forward
set forward-slot master
set priority 5
set comment "ipv6 multicast"
end
```

Controlling SNAT port partitioning behavior

When you set up FortiGate 6000F SNAT configurations using IP pools, the source NAT (SNAT) source ports in the IP pool are distributed evenly among the FPCs. Each FPC gets an equal share of the source ports. For details about how this works and issues that occur, see the Fortinet Community article [Technical Tip: Understanding NAT port allocation on Chassis \(6k/7k\)](#).

You can use the following command to control how the FortiGate 6000F partitions SNAT source ports among FPCs:

```
config load-balance setting
set nat-source-port {chassis-slots | enabled-slots}
end
```

`chassis-slots` this option statically allocates SNAT source ports to all FPCs that are enabled when you enter the command. If you disable an FPC from the CLI, the SNAT source ports assigned to that FPC will not be re-allocated to the remaining FPCs. All FPCs that are still operating will maintain the same SNAT source port allocation and active sessions being processed by the still operating FPCs will not be affected.



You can use the following command to enable or disable an FPC from the CLI:

```
config workers
  edit <slot>
    set status {disable | enable}
  end
```

`enabled-slots` this option dynamically re-distributes SNAT source ports to enabled FPCs. This is the default behavior and is recommended in most cases.

If an FPC is disabled, SLBC dynamically re-allocates SNAT source ports among the remaining enabled FPCs. This means that all configured SNAT source ports remain available. If SNAT source ports are re-allocated when the FortiGate 6000F is actively processing traffic, some active sessions may be lost if their source ports are allocated to different FPCs.



SNAT source ports are not dynamically reallocated if an FPC is powered off. To re-allocate SNAT source ports, the FPC must be disabled from the CLI.

Showing how the DP3 processor will load balance a session

You can use the following command to display the FPC slot that the DP3 processor will load balance a session to.

```
diagnose load-balance dp find session {normal | reverse | fragment | pinhole}
```

Normal and reverse sessions

For a normal or corresponding reverse session you can define the following:

```
{normal | reverse} <ip-protocol> <src-ip> {<src-port> | <icmp-type> | <icmp-typecode>} <dst-  
ip> {<dst-port> | <icmp-id>} [<x-vid>] [<x-cfi>] [<x-pri>]
```

Fragment packet sessions

For a session for fragment packets you can define the following:

```
fragment <ip-protocol> {<src-port> | <icmp-type> | <icmp-typecode>} <dst-ip> <ip-id> [<x-  
vid>] [<x-cfi>] [<x-pri>]
```

Pinhole sessions

For a pinhole sessions you can define the following:

```
pinhole <ip-protocol> <dst-ip> <dst-port> [<x-vid>] [<x-cfi>] [<x-pri>]
```

Normal session example output

For example, the following command shows that a new TCP session (protocol number 6) with source IP address 11.1.1.11, source port 53386, destination IP address 12.1.1.11, and destination port 22 would be sent to TCP slot 8 by the DP3 processor.

```
diagnose load-balance dp find session normal 6 11.1.1.11 53386 12.1.1.11 22
=====
MBD SN: F6KF503E17900068
Primary Bin 9708928
New session to slot 8 (src-dst-ip-sport-dport)
```

Additional information about the session also appears in the command output in some cases.

Maximum number of flow rules limited by hardware

For all FortiGate-6000 models, the CLI allows you to add up to 512 flow rules. However, the number of flow rules that you can add is actually limited by the FortiGate-6000 internal switch hardware and all FortiGate-6000 models support up to 256 flow rules.

SSL VPN load balancing

FortiGate-6000 supports load balancing SSL VPN tunnel mode sessions terminated by the FortiGate-6000. By default SSL VPN load balancing is disabled and a flow rule is required to send all SSL VPN sessions to one FPC (usually the primary FPC).

To support SSL VPN tunnel load balancing, you must disable all flow rules that match the SSL VPN traffic to be load balanced.

For SSL VPN load balancing to work properly, the DP processor load distribution method must be changed to a setting that does not include `src-port`. The following DP load distribution methods are supported for SSL VPN load balancing:

```
config load balance setting
    set dp-load-distribution-method {to-master | src-ip | dist-ip | src-dst-ip | dis-ip-
        dport}
end
```

Then you can use the following command to enable SSL VPN load balancing:

```
config load-balance setting
    set sslvpn-load-balance enable
end
```

When you enable SSL VPN load balancing, the FortiGate-6000 restarts SSL VPN processes running on the management board and the FPCs, resetting all current SSL VPN sessions. This restart will interrupt any active SSL VPN sessions.

Once the SSL VPN processes restart, the FortiGate-6000 DP3 processor distributes SSL VPN tunnel mode sessions to all of the FPCs.

To be able to distribute SSL VPN sessions to all FPCs, SSL VPN load balancing statically allocates the IP addresses in SSL VPN IP pools among the FPCs. Each FPC acquires a subset of the IP addresses in the IP pool. You may need to expand the number of IP addresses in your SSL VPN IP pools to make sure enough IP addresses are available for each FPC.



SSL VPN IP pool IP addresses are not re-allocated if an FPC goes down, is disabled, or is taken offline. The IP pool IP addresses assigned to the missing FPC are not available until the FPC returns to normal operation.

No other special configuration is required to support SSL VPN tunnel mode load balancing.

For more information on FortiGate 6000F SSL VPN load balancing, see this Fortinet Community article: [Technical Tip : How to load balance SSL VPN web-mode traffic on FortiGate-6000 series.](#)

Setting up SSL VPN using flow rules

As an alternative to SSL VPN load balancing, you can manually add SSL VPN load balancing flow rules to configure the FortiGate-6000 to send all SSL VPN sessions to the primary FPC. To match SSL VPN traffic, the flow rule should include a destination port that matches the destination port of the SSL VPN server. A basic rule to send SSL VPN traffic to the primary FPC could be:

```

config load-balance flow-rule
  edit 0
    set status enable
    set ether-type ipv4
    set protocol tcp
    set dst-l4port 443-443
    set forward-slot master
    set comment "ssl vpn server to primary worker"
  end

```

This flow rule matches all sessions sent to port 443 (the default SSL VPN server listening port) and sends these sessions to the primary FPC. This should match all of your SSL VPN traffic if you are using the default SSL VPN server listening port (443). This flow rule also matches all other sessions using 443 as the destination port so all of this traffic is also sent to the primary FPC.



If you want to set up SSL VPN using flow rules, you should use the following command to make sure SSL VPN load balancing is disabled:

```

config load-balance setting
  set sslvpn-load-balance disable
end

```

Also, there are no restrictions on the DP load balancing method if you are setting up SSL VPN using flow rules.

If you change the SSL VPN server listening port

As a best practice, if you add a flow rule for SSL VPN, Fortinet recommends using a custom SSL VPN port (for example, 10443 instead of 443). This can improve performance by allowing SSL traffic on port 443 that is not part of your SSL VPN to be load balanced to FPCs instead of being sent to the primary FPC by the SSL VPN flow rule.

If you have changed the SSL VPN server listening port to 10443, you can change the SSL VPN flow rule as follows:

```

config load-balance flow-rule
  edit 26
    set status enable
    set ether-type ipv4
    set protocol tcp
    set dst-l4port 10443-10443
    set forward-slot master
    set comment "ssl vpn server to primary worker"
  end

```

You can also make the SSL VPN flow rule more specific by including the SSL VPN server interface in the flow rule. For example, if your FortiGate-6000 listens for SSL VPN sessions on the port12 interface:

```

config load-balance flow-rule
  edit 26
    set status enable
    set ether-type ipv4
    set protocol tcp
    set src-interface port12
    set dst-l4port 10443-10443
    set forward-slot master
    set comment "ssl vpn server to primary worker"
  end

```

Adding the SSL VPN server IP address

You can also add the IP address of the FortiGate-6000 interface that receives SSL VPN traffic to the SSL VPN flow rule to make sure that the flow rule only matches the traffic of SSL VPN clients connecting to the SSL VPN server. For example, if the IP address of the interface is 172.25.176.32:

```
config load-balance flow-rule
  edit 26
    set status enable
    set ether-type ipv4
    set protocol tcp
    set dst-addr-ipv4 172.25.176.32 255.255.255.255
    set dst-l4port 10443-10443
    set forward-slot master
    set comment "ssl vpn server to primary worker"
  end
```

This flow rule will now only match SSL VPN sessions with 172.25.176.32 as the destination address and send all of these sessions to the primary FPC.

IPsec VPN load balancing

The FortiGate 6000F uses SLBC load balancing to select an FPC to terminate traffic for a new IPsec VPN tunnel instance and all traffic for that tunnel instance is terminated on the same FPC.

```
config vpn ipsec phase1-interface
  edit <name>
    set ipsec-tunnel-slot {auto | FPC1 | FPC2 | FPC3 | FPC4 | FPC5 | FPC6 | FPC7 | FPC8 |
      FPC9 | FPC10 | master}
  end
```

You can optionally use the IPsec tunnel phase 1 configuration to select a specific FPC to terminate all tunnel instances started by that phase 1. For example, to terminate all tunnels on FPC5:

```
config vpn ipsec phase1-interface
  edit <name>
    set ipsec-tunnel-slot FPC5
  end
```

FortiGate 6000F IPsec VPN supports the following features:

- Interface-based IPsec VPN (also called route-based IPsec VPN).
- Site-to-Site IPsec VPN.
- Dialup IPsec VPN. The FortiGate 6000F can be the dialup server or client.
- Static and dynamic routing (BGP, OSPF, and RIP) over IPsec VPN tunnels.
- When an IPsec VPN tunnel is initialized, the SA is synchronized to all FPCs in the FortiGate 6000F, or in both FortiGate 6000Fs in an HA configuration.
- Traffic between IPsec VPN tunnels is supported when both tunnels terminate on the same FPC.
- When setting up a VRF configuration to send traffic between two IPsec VPN interfaces with different VRFs, both IPsec tunnels must terminate on the same FPC.

FortiGate 6000F IPsec VPN has the following limitations:

- Policy-based IPsec VPN tunnels terminated by the FortiGate 6000F are not supported.
- Policy routes cannot be used for communication over IPsec VPN tunnels.
- IPv6 clear text traffic over IPv4 or IPv6 IPsec tunnels terminated on the FortiGate 6000F is not supported. This limitation does not affect IPv4 clear text traffic over IPv4 or IPv6 IPsec tunnels terminated on the FortiGate 6000F. This limitation also does not affect any pass through IPsec tunnel traffic that does not terminate on the FortiGate 6000F.
- The FortiGate 6000F, because it uses DP processors for SLBC, does not support IPsec VPN to remote networks with 0- to 15-bit netmasks.

Configuring IPsec VPN load balancing

FortiGate 6000F IPsec load balancing is tunnel based. You can set the load balance strategy for each tunnel when configuring `phase1-interface` options:

```
config vpn ipsec phase1-interface
  edit <name>
```

```

    set ipsec-tunnel-slot {auto | FPC1 | FPC2 | FPC3 | FPC4 | FPC5 | FPC6 | FPC7 | FPC8 |
        FPC9 | FPC10 | master}
end

```

`auto` the default setting. All tunnels started by this phase 1 are load balanced to an FPC slot based on the `src-ip` and `dst-ip` hash result. All traffic for a given tunnel instance is processed by the same FPC.

`FPC1` to `FPC10` all tunnels started by this phase 1 terminate on the selected FPC. For the FortiGate-6300F and 6301F the options are `FPC1` to `FPC6`.

`master` all tunnels started by this phase 1 terminate on the primary FPC.

Even if you select `master` or a specific FPC, new SAs created by this tunnel are synchronized to all FPCs.

SD-WAN with multiple IPsec VPN tunnels

To support SD-WAN with IPsec VPN, the IPsec VPN tunnel configuration of all IPsec VPN tunnels that are members of the same SD-WAN zone in the same VDOM must send traffic to the same FPC. This means the `ipsec-tunnel-slot` configuration of the IPsec VPN tunnel must include a specific FPC. Setting `ipsec-tunnel-slot` to `master` is not recommended, since the primary FPC can change. Setting `ipsec-tunnel-slot` to `auto` is not supported.

SD-WAN with multiple IPsec VPN tunnels on a FortiGate 6000F has the following limitations:

- Auto negotiation must be enabled in the IPsec VPN phase 2 configuration for all IPsec tunnels added to an SD-WAN zone.
- An SD-WAN zone can include a mixture of IPsec VPN interfaces and other interface types (for example, physical interfaces). If an SD-WAN zone contains an IPsec VPN interface, all traffic accepted by interfaces in that SD-WAN zone is sent to the same FPC, including traffic accepted by other interface types.
- SD-WAN health checking is not supported for IPsec VPN SD-WAN members.
- SD-WAN traffic information, including packet statistics, policy hit counts, and so on is not supported for IPsec VPN SD-WAN members.

Example FortiGate 6000F IPsec VPN VRF configuration

The following shows the basics of how to set up a VRF configuration that allows traffic between two IPsec VPN interfaces with different VRFs on a FortiGate 6000F. To support this configuration, both IPsec tunnels must terminate on the same FPC, in this example, the FPC in slot 5.

Create two VLAN interfaces:

```

config system interface
  edit "v0031"
    set vdom "vrf1"
    set vrf 10
    set ip <ip-address>
    set interface "port1"
    set vlanid 31
  next
  edit "v0032"
    set vdom "vrf1"
    set vrf 11

```

```

    set ip <ip-address>
    set interface "port2"
    set vlanid 32
next

```

Create two phase1-interface tunnels. Add each tunnel to one of the VLAN interfaces created in step 1. The `ipsec-tunnel-slot` setting for both is FPC5.

```

config vpn ipsec phase1-interface
  edit "p1-v31"
    set interface "v0031"
    set local-gw <ip-address>
    set peertype any
    set proposal 3des-sha256
    set remote-gw <ip-address>
    set psksecret <psk>
    set ipsec-tunnel-slot FPC5
  next
  edit "p1-v32"
    set interface "v0032"
    set local-gw <ip-address>
    set peertype any
    set proposal 3des-sha256
    set remote-gw <ip-address>
    set psksecret <psk>
    set ipsec-tunnel-slot FPC5
end

```

Edit each IPsec VPN interface and set the VRF ID for each one:

```

config system interface
  edit "p1-v31"
    set vdom "vrf1"
    set vrf 10
    set type tunnel
    set interface "v0031"
  next
  edit "p1-v32"
    set vdom "vrf1"
    set vrf 11
    set type tunnel
    set interface "v0032"
end

```

Troubleshooting

Use the following commands to verify that IPsec VPN sessions are up and running.

Use the `diagnose load-balance status` command from the management board to determine the primary FPC. The third line of the command output shows which FPC is operating as the primary FPC.

```

diagnose load-balance status
=====
MBD SN: F6KF313E17900032
  Primary FPC Blade: slot-5

```

```

Slot 1: FPC6KF3E17900200
  Status:Working  Function:Active
  Link:          Base: Up          Fabric: Up
  Heartbeat: Management: Good    Data: Good
  Status Message:"Running"
Slot 2: FPC6KF3E17900201
  Status:Working  Function:Active
  Link:          Base: Up          Fabric: Up
  Heartbeat: Management: Good    Data: Good
  Status Message:"Running"
Slot 3: FPC6KF3E17900207
  Status:Working  Function:Active
  Link:          Base: Up          Fabric: Up
  Heartbeat: Management: Good    Data: Good
  Status Message:"Running"
Slot 4: FPC6KF3E17900219
  Status:Working  Function:Active
  Link:          Base: Up          Fabric: Up
  Heartbeat: Management: Good    Data: Good
  Status Message:"Running"
Slot 5: FPC6KF3E17900235
  Status:Working  Function:Active
  Link:          Base: Up          Fabric: Up
  Heartbeat: Management: Good    Data: Good
  Status Message:"Running"
Slot 6: FPC6KF3E17900169
  Status:Working  Function:Active
  Link:          Base: Up          Fabric: Up
  Heartbeat: Management: Good    Data: Good
  Status Message:"Running"

```

Log into the primary FPC CLI and run the command `diagnose vpn tunnel list <phase2-name>` to show the sessions for the phase 2 configuration. The example below is for the `to-fgt2` phase 2 configuration configured previously in this chapter. The command output shows the security association (SA) setup for this phase 2 and the all of the destination subnets .

From the command output, make sure the SA is installed and the `dst` addresses are correct.

```

diagnose vpn tunnel list name to-fgt2
list ipsec tunnel by names in vd 11
-----
name=to-fgt2 ver=1 serial=2 4.2.0.1:0->4.2.0.2:0
bound_if=199 lgwy=static/1 tun=intf/0 mode=auto/1 encap=none/40 options[0028]=npu ike_assit
proxyid_num=1 child_num=0 refcnt=8581 ilast=0 olast=0 auto-discovery=0
ike_asssit_last_sent=4318202512
stat: rxp=142020528 txp=147843214 rxb=16537003048 txb=11392723577
dpd: mode=on-demand on=1 idle=20000ms retry=3 count=0 seqno=2
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=to-fgt2 proto=0 sa=1 ref=8560 serial=8
  src: 0:4.2.1.0/255.255.255.0:0 0:4.2.2.0/255.255.255.0:0
  dst: 0:4.2.3.0/255.255.255.0:0 0:4.2.4.0/255.255.255.0:0 0:4.2.5.0/255.255.255.0:0
  SA: ref=7 options=22e type=00 soft=0 mtu=9134 expire=42819/0B replaywin=2048 seqno=4a26f
  esn=0 replaywin_lastseq=00045e80
  life: type=01 bytes=0/0 timeout=43148/43200
  dec: spi=e89caf36 esp=aes key=16 26aa75c19207d423d14fd6fef2de3bcf
      ah=sha1 key=20 7d1a330af33fa914c45b80c1c96eafaf2d263ce7

```

```

enc: spi=b721b907 esp=aes key=16 acb75d21c74eabc58f52ba96ee95587f
ah=sha1 key=20 41120083d27eb1d3c5c5e464d0a36f27b78a0f5a
dec:pkts/bytes=286338/40910978, enc:pkts/bytes=562327/62082855
npu_flag=03 npu_rgwy=4.2.0.2 npu_lgwy=4.2.0.1 npu_selid=b dec_npuid=3 enc_npuid=1

```

Log into the CLI of any of the FPCs and run the command `diagnose test application fctrlproxyd 2`. The output should show matching destination subnets.

```
diagnose test application fctrlproxyd 2
```

```
fcp route dump : last_update_time 24107
```

```
Slot:4
```

```

routecache entry: (5)
checksum:27 AE 00 EA 10 8D 22 0C D6 48 AB 2E 7E 83 9D 24
vd:3 p1:to-fgt2 p2:to-fgt2 subnet:4.2.3.0 mask:255.255.255.0 enable:1
vd:3 p1:to-fgt2 p2:to-fgt2 subnet:4.2.4.0 mask:255.255.255.0 enable:1
vd:3 p1:to-fgt2 p2:to-fgt2 subnet:4.2.5.0 mask:255.255.255.0 enable:1
=====

```

FortiGate-6000 high availability

FortiGate-6000 for FortiOS 6.4 supports the following types of HA operation:

- FortiGate Clustering Protocol (FGCP)
- Virtual clustering ([Virtual clustering on page 68](#))
- FortiGate Session Life Support Protocol (FGSP) ([FortiGate-6000 FGSP on page 79](#))
- Inter-cluster session synchronization ([Synchronizing sessions between FGCP clusters on page 82](#))
- Standalone configuration synchronization ([Standalone configuration synchronization on page 87](#))
- Virtual Router Redundancy Protocol (VRRP) ([FortiGate-6000 VRRP HA on page 89](#))

Introduction to FortiGate-6000 FGCP HA

FortiGate-6000 supports active-passive FortiGate Clustering Protocol (FGCP) HA between two (and only two) identical FortiGate-6000s. You can configure FortiGate-6000 HA in much the same way as any FortiGate HA setup except that only active-passive HA is supported.

You must use the 10Gbit HA1 and HA2 interfaces for HA heartbeat communication. The recommended HA heartbeat configuration is to use a cable to directly connect the HA1 interfaces of each FortiGate-6000 and another cable to directly connect the HA2 interfaces of each FortiGate-6000.

You can use switches to connect the HA heartbeat interfaces. Heartbeat packets are VLAN-tagged and you can configure the VLANs used. If you are using switches you must configure the switch interfaces in trunk mode and the switches must allow the VLAN-tagged packets.

As part of the FortiGate-6000 HA configuration, you assign each of the FortiGate-6000s in the HA cluster a chassis ID of 1 or 2. The chassis IDs just allow you to identify individual FortiGate-6000s and do not influence primary unit selection.

If both FortiGate-6000s in a cluster are configured with the same chassis ID, both chassis begin operating in HA mode without forming a cluster. A message similar to the following is displayed on the CLI console of both devices:

```
HA cannot be formed because this box's chassis-id 1 is the same from the  
HA peer 'F76E9D3E17000001' chassis-id 1.
```

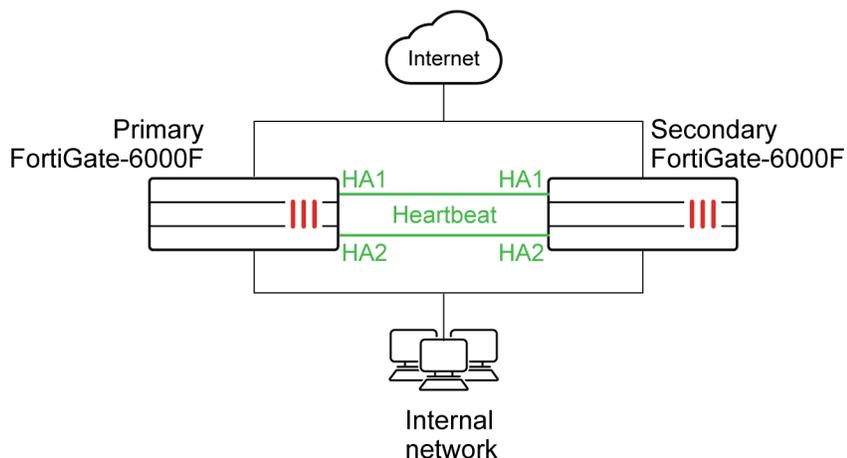


As well, a log message similar to the following is created:

```
Jan 29 16:29:46 10.160.45.70 date=2020-01-29 time=16:29:51 devname="CH-  
02" devid="F76E9D3E17000001" slot=1 logid="0108037904" type="event"  
subtype="ha" level="error" vd="mgmt-vdom" eventtime=1580344192162305962  
tz="-0800" logdesc="Device set as HA primary" msg="HA group detected  
chassis-id conflict" ha_group=7 sn="F76E9DT018900001 chassis-id=1"
```

You can resolve this issue by logging into one of the FortiGates and changing its Chassis ID to 2. When this happens, the two chassis will form a cluster.

Example FortiGate-6000 HA configuration



In a FortiGate-6000 FGCP HA configuration, the primary FortiGate-6000 processes all traffic. The secondary FortiGate-6000 operates in hot standby mode. The FGCP synchronizes the configuration, active sessions, routing information, and so on to the secondary FortiGate-6000. If the primary FortiGate-6000 fails, traffic automatically fails over to the secondary.

Before you begin configuring HA

Before you begin:

- The FortiGate-6000s must be running the same FortiOS firmware version
- Interfaces should be configured with static IP addresses (not DHCP or PPPoE).
- Register and apply licenses to each FortiGate-6000 before setting up the HA cluster. This includes licensing for FortiCare, IPS, AntiVirus, Web Filtering, Mobile Malware, FortiClient, FortiCloud, and additional virtual domains (VDMs).
- Both FortiGate-6000s in the cluster must have the same level of licensing for FortiGuard, FortiCloud, FortiClient, and VDMs.
- FortiToken licenses can be added at any time because they are synchronized to all cluster members.
- Both FortiGate 6001Fs, FortiGate-6501Fs or FortiGate-6301Fs in a cluster must have the same number of active hard disks and the same RAID configuration. Use the `execute disk list` command to confirm the log disk and RAID configuration of each device.

On each FortiGate-6000, make sure the configurations of the FPCs are synchronized before starting to configure HA. You can use the following command to verify the configuration status of the FPCs. The following example shows the results for a FortiGate-6300F.

```
diagnose sys confsync showchsum | grep all
all: c0 68 d2 67 e1 23 d9 3a 10 50 45 c5 50 f1 e6 8e
all: c0 68 d2 67 e1 23 d9 3a 10 50 45 c5 50 f1 e6 8e
all: c0 68 d2 67 e1 23 d9 3a 10 50 45 c5 50 f1 e6 8e
all: c0 68 d2 67 e1 23 d9 3a 10 50 45 c5 50 f1 e6 8e
all: c0 68 d2 67 e1 23 d9 3a 10 50 45 c5 50 f1 e6 8e
all: c0 68 d2 67 e1 23 d9 3a 10 50 45 c5 50 f1 e6 8e
all: c0 68 d2 67 e1 23 d9 3a 10 50 45 c5 50 f1 e6 8e
```

If the FPCs are synchronized, the listed checksums should all be the same.

You can also use the following command to list the FPCs that are synchronized. The example output, for a FortiGate-6300F, shows all six FPCs have been configured for HA and added to the cluster.

```
diagnose sys confsync status | grep in_sync
F6KF313E17900031, Secondary, uptime=232441.23, priority=2, slot_id=1:0, idx=0, flag=0x10, in_sync=1
F6KF313E17900031, Primary, uptime=232441.23, priority=1, slot_id=1:0, idx=0, flag=0x10, in_sync=1
F6KF313E17900031, Secondary, uptime=232441.23, priority=2, slot_id=1:0, idx=0, flag=0x10, in_sync=1
FPC6KF3E17900209, Secondary, uptime=231561.99, priority=24, slot_id=1:6, idx=6, flag=0x24, in_sync=1
FPC6KF3E17900215, Secondary, uptime=231524.81, priority=22, slot_id=1:4, idx=7, flag=0x24, in_sync=1
FPC6KF3E17900217, Secondary, uptime=232289.83, priority=120, slot_id=1:5, idx=8, flag=0x24, in_sync=1
FPC6KF3E17900229, Secondary, uptime=232271.59, priority=118, slot_id=1:3, idx=10, flag=0x24, in_sync=1
FPC6KF3E17900230, Secondary, uptime=232330.19, priority=116, slot_id=1:1, idx=11, flag=0x24, in_sync=1
FPC6KF3E17900291, Secondary, uptime=232314.29, priority=117, slot_id=1:2, idx=13, flag=0x24, in_sync=1
```

In this command output `in_sync=1` means the FPC is synchronized with the management board and `in_sync=0` means the FPC is not synchronized.

Connect the HA1 and HA2 interfaces for HA heartbeat communication

HA heartbeat communication between FortiGate-6000s happens over the 10Gbit HA1 and HA2 interfaces. To set up HA heartbeat connections:

- Connect the HA1 interfaces of the two FortiGate-6000s together either with a direct cable connection, or using a switch.
- Connect the HA2 interfaces in the same way.

Using separate connections for HA1 and HA2 is recommended for redundancy. Direct cable connections between the HA heartbeat interfaces is also recommended.

If you are using switches to connect the HA1 and HA2 interfaces, the switches need to be configured in trunk mode. It is also recommended that these switches be dedicated to HA heartbeat communication and not used for other traffic. But

You can use the same switch for both HA1 and HA2 as long as you separate the HA1 and HA2 traffic on the switch. To do this, enable trunk mode for the switch interfaces, and set the heartbeat traffic on the HA1 and HA2 Interfaces to have different VLAN IDs. See the following sections for information about the configuration options you can use to customize HA heartbeat packets to be compatible with different third-party switch configurations.



FortiGate-6000 HA does not support using the HA `session-sync-dev` option. Instead, session synchronization traffic uses the HA1 and HA2 interfaces, separating session sync traffic from data traffic

Default HA heartbeat VLAN triple-tagging

By default, HA heartbeat packets are VLAN packets with VLAN ID 999, an outer TPID of 0x8100, ethertype 8890, and an MTU value of 1500. The default proprietary HA heartbeat VLAN tagging uses the following triple tagging format:

TPID 0x8100 VLAN <vlan-id> (by default 999) + TPID 0x88a8 VLAN 10/30 + TPID 0x8100 VLAN 10/30 + ethernet packet

If your switch is compatible with Fortinet's proprietary triple-tagging format then all you need to do is use the following options to give the M1 and M2 interfaces different VLAN tags.

```
config system ha
  set ha-port-dtag-mode proprietary
  set hbdev-vlan-id <vlan>
  set hbdev-second-vlan-id <vlan>
end
```

Where:

- `ha-port-dtag-mode` is set to `proprietary` and the FortiGate-6000 uses the default triple-tagging format.
- `hbdev-vlan-id` sets the outer VLAN ID used by HA1 interface heartbeat packets.
- `hbdev-second-vlan-id` sets the outer VLAN ID used by HA2 interface heartbeat packets. The HA1 and HA2 interfaces must have different outer VLAN IDs if they are connected to the same switch.

If your switch is not compatible with Fortinet's proprietary triple-tagging format, you can use the following options to change the outer TPID and ethertype.

```
config system ha
  set ha-port-dtag-mode proprietary
  set ha-port-outer-tpid {0x8100 | 0x88a8 | 0x9100}
  set ha-eth-type <ethertype>
end
```

Where:

- `ha-port-dtag-mode` is set to `proprietary` and the FortiGate-6000 uses the default triple-tagging format.
- `ha-port-outer-tpid` sets the outer TPID to be compatible with the switch. The default outer TPID of 0x8100, is compatible with most third-party switches.
- `ha-eth-type` sets the HA heartbeat packet ethertype (default 8890) to be compatible with the switch.



If your switch doesn't support triple tagging, see [HA heartbeat VLAN double-tagging on page 56](#).

Example triple-tagging compatible switch configuration

The switch that you use for connecting HA heartbeat interfaces does not have to support IEEE 802.1ad (also known as Q-in-Q, double-tagging). But the switch should be able to forward the double-tagged frames. Some switches will strip out the inner tag and Fortinet recommends avoiding these switches. FortiSwitch D and E series can correctly forward double-tagged frames.



This configuration is not required for FortiGate-6000 HA configurations if you have set up direct connections between the HA heartbeat interfaces.

This example shows how to configure a FortiGate-6000 to use different VLAN IDs for the HA1 and HA2 HA heartbeat interfaces and then how to configure two interfaces on a Cisco switch to allow HA heartbeat packets.



This example sets the native VLAN ID for both switch ports to 777. You can use any VLAN ID as the native VLAN ID as long as the native VLAN ID is not the same as the allowed VLAN ID.

1. On both FortiGate-6000s, enter the following command to use different VLAN IDs for the HA1 and HA2 interfaces. The command sets the ha1 VLAN ID to 4091 and the ha2 VLAN ID to 4092:

```
config system ha
  set ha-port-dtag-mode proprietary
  set hbdev ha1 50 ha2 100
  set hbdev-vlan-id 4091
  set hbdev-second-vlan-id 4092
end
```

2. Use the `get system ha` or `get system ha status` command to confirm the VLAN IDs.

```
get system ha status
...
HBDEV stats:
F6KF51T018900026(updated 4 seconds ago):
  ha1: physical/10000full, up, rx-bytes/packets/dropped/errors=54995955/230020/0/0,
tx=63988049/225267/0/0, vlan-id=4091
  ha2: physical/10000full, up, rx-bytes/packets/dropped/errors=54995955/230020/0/0,
tx=63988021/225267/0/0, vlan-id=4092
F6KF51T018900022(updated 3 seconds ago):
  ha1: physical/10000full, up, rx-bytes/packets/dropped/errors=61237440/230023/0/0,
tx=57746989/225271/0/0, vlan-id=4091
  ha2: physical/10000full, up, rx-bytes/packets/dropped/errors=61238907/230023/0/0,
tx=57746989/225271/0/0, vlan-id=4092
...
```

3. Configure the Cisco switch interface that connects the HA1 interfaces to allow packets with a VLAN ID of 4091:

```
interface <name>
switchport mode trunk
switchport trunk native vlan 777
switchport trunk allowed vlan 4091
```

4. Configure the Cisco switch port that connects the HA2 interfaces to allow packets with a VLAN ID of 4092:

```
interface <name>
switchport mode trunk
switchport trunk native vlan 777
switchport trunk allowed vlan 4092
```

HA heartbeat VLAN double-tagging

FortiGate-6000 HA supports HA heartbeat double-tagging to be compatible with third-party switches that do not support Fortinet's proprietary triple tagging format. HA heartbeat double-tagging has the following format:

TPID 0x8100 VLAN <vlan-id> (by default 999) + TPID 0x8100 VLAN 10/30 + ethernet packet

You can use the following commands to set the HA VLAN tagging mode to double-tagging, customize the outer TPID, and set the VLAN IDs for HA1 and HA2. Both FortiGates in the cluster must have the same VLAN tagging configuration.

```
config system ha
  set ha-port-dtag-mode double-tagging
  set ha-port-outer-tpid {0x8100 | 0x88a8 | 0x9100}
  set hbdev-vlan-id <vlan>
```

```

    set hbdev-second-vlan-id <vlan>
    set ha-eth-type <ethertype>
end

```

Where:

`ha-port-dtag-mode` is set to `double-tagging` and the FortiGate-6000 uses the double-tagging format.

`ha-port-outer-tpid` sets the outer TPID to be compatible with the switch. The default outer TPID of 0x8100 is compatible with most third-party switches.

`hbdev-vlan-id` sets the outer VLAN ID used by HA1 interface heartbeat packets.

`hbdev-second-vlan-id` sets the outer VLAN ID used by HA2 interface heartbeat packets. The HA1 and HA2 interfaces must have different outer VLAN IDs if they are connected to the same switch.

`ha-eth-type` sets the HA heartbeat packet ethertype (default 8890) to be compatible with the switch.

Example double-tagging switch configuration

The following switch configuration is compatible with FortiGate-6000 HA heartbeat double tagging and with the default TPID of 0x8100.

The FortiGate-6000 HA heartbeat configuration is.

```

config system ha
    set ha-port-dtag-mode double-tagging
    set hbdev ha1 50 ha2 50
    set hbdev-vlan-id 4091
    set hbdev-second-vlan-id 4092
end

```

Example third-party switch configuration:

Switch interfaces 37 and 38 connect to the HA1 interfaces of both FortiGate-6000s.

```

interface Ethernet37
description **** FGT-6000F HA1 HA HB ****
speed forced 10000full
switchport access vlan 660
switchport trunk native vlan 4091
switchport mode dot1q-tunnel
!
interface Ethernet38
description **** FGT-6000F HA1 HA HB ****
speed forced 10000full
switchport access vlan 660
switchport trunk native vlan 4091
switchport mode dot1q-tunnel
!

```

Switch interfaces 39 and 40 connect to the HA2 interfaces of both FortiGate-6000s.

```

interface Ethernet39
description **** FGT-6000F HA2 HA HB ****
mtu 9214
speed forced 10000full
no error-correction encoding

```

```

switchport access vlan 770
switchport trunk native vlan 4092
switchport mode dot1q-tunnel
!
interface Ethernet42
description **** FGT-6000F HA2 HA HB ****
mtu 9214
speed forced 10000full
no error-correction encoding
switchport access vlan 770
switchport trunk native vlan 4092
switchport mode dot1q-tunnel
!

```

Basic FortiGate-6000 HA configuration

Use the following steps to set up HA between two FortiGate-6000s. To configure HA, you assign a chassis ID (1 and 2) to each of the FortiGate-6000s. These IDs allow the FGCP to identify the chassis and do not influence primary FortiGate selection. Before you start, determine which FortiGate-6000 should be chassis 1 and which should be chassis 2.

Make sure you give each FortiGate-6000 a different chassis ID. If both FortiGate-6000s in a cluster are configured with the same chassis ID, both chassis begin operating in HA mode without forming a cluster. A message similar to the following is displayed on the CLI console of both devices:

```

HA cannot be formed because this box's chassis-id 1 is the same from the
HA peer 'F76E9D3E17000001' chassis-id 1.

```

As well, a log message similar to the following is created:

```

Jan 29 16:29:46 10.160.45.70 date=2020-01-29 time=16:29:51 devname="CH-
02" devid="F76E9D3E17000001" slot=1 logid="0108037904" type="event"
subtype="ha" level="error" vd="mgmt-vdom" eventtime=1580344192162305962
tz="-0800" logdesc="Device set as HA primary" msg="HA group detected
chassis-id conflict" ha_group=7 sn="F76E9DT018900001 chassis-id=1"

```



You can resolve this issue by logging into one of the FortiGate-6000s and changing its Chassis ID to 2. When this happens, the two chassis will form a cluster.

Also, if you are setting up a cluster of FortiGate-6301Fs or 6501Fs, before you configure HA, consider using the `execute disk list` command on each FortiGate to verify that they both have the same disk and RAID configuration. If one of the FortiGates only has one operating hard disk, when the cluster forms the FortiGate with fewer operating hard disks will be shut down. If the RAID configurations of each FortiGate are different, when the cluster forms, the FortiGate with the lowest RAID level will be shut down. You can use the `execute disk format` command to format the disks and the `execute disk raid` command to set both FortiGates to the same RAID mode.

1. Set up HA heartbeat communication as described in [Connect the HA1 and HA2 interfaces for HA heartbeat communication on page 54](#).
2. Log into the GUI or CLI of the FortiGate-6000 that will become chassis 1.

- Use the following CLI command to change the host name. This step is optional, but setting a host name makes the FortiGate-6000 easier to identify after the cluster has formed.

```
config system global
  set hostname 6K-Chassis-1
end
```

From the GUI you can configure the host name by going to **System > Settings** and changing the **Host name**.

- Enter the following command to configure basic HA settings for the chassis 1 FortiGate-6000.

```
config system ha
  set group-id 6
  set group-name My-6K-cluster
  set mode a-p
  set hbdev ha1 50 ha2 100
  set chassis-id 1
  set password <password>
end
```

From the GUI you can configure HA by going to **System > HA**. Set the **Mode** to **Active-Passive**, set the **Group Name**, add a **Password**, select the **Chassis identifier** (or chassis ID), and set the **Heartbeat Interface Priority** for the heartbeat interfaces (HA1 and HA2). You must configure the group ID from the CLI.

- If you are connecting the HA heartbeat interfaces together with a switch, change the HA heartbeat VLAN IDs, for example:

```
config system ha
  set hbdev-vlan-id 4091
  set hbdev-second-vlan-id 4092
end
```

- Log into the chassis 2 FortiGate-6000 and configure its host name, for example:

```
config system global
  set hostname 6K-Chassis-2
end
```

From the GUI you can configure the host name by going to **System > Settings** and changing the **Host name**.

- Enter the following command to configure basic HA settings. The configuration must be the same as the chassis 1 configuration, except for the chassis ID.

```
config system ha
  set group-id 6
  set group-name My-6K-cluster
  set mode a-p
  set hbdev ha1 50 ha2 100
  set chassis-id 2
  set password <password>
end
```

From the GUI you can configure HA by going to **System > HA**. Set the **Mode** to **Active-Passive**, set the **Group Name**, add a **Password**, select the **Chassis identifier**, and set the **Heartbeat Interface Priority** for the heartbeat interfaces (HA1 and HA2). You must configure the group ID from the CLI.

- If you are connecting the HA heartbeat interfaces together with a switch, change the HA heartbeat VLAN IDs, for example:

```
config system ha
  set hbdev-vlan-id 4091
  set hbdev-second-vlan-id 4092
end
```

Once you save your configuration changes, if the HA heartbeat interfaces are connected, the FortiGate-6000s negotiate to establish a cluster. You may temporarily lose connectivity with the FortiGate-6000s as the cluster negotiates and the FGCP changes the MAC addresses of the FortiGate-6000 interfaces.

9. Log into the cluster and view the HA Status dashboard widget or enter the `get system ha status` command to confirm that the cluster has formed and is operating normally.

If the cluster is operating normally, you can connect network equipment, add your configuration, and start operating the cluster.

Verifying that the cluster is operating normally

You view the cluster status from the HA Status dashboard widget, by going to **System > HA**, or by using the `get system ha status` command.

If the HA Status widget or the `get system ha status` command shows a cluster has not formed, check the HA heartbeat connections. They should be configured as described in [Connect the HA1 and HA2 interfaces for HA heartbeat communication on page 54](#).

You should also review the HA configurations of the FortiGate-6000s. When checking the configurations, make sure both FortiGate-6000s have the same HA configuration, including identical HA group IDs, group names, passwords, and HA heartbeat VLAN IDs. Also make sure the FortiGate-6000s have different chassis IDs.

The following example FortiGate-6000 `get system ha status` output shows a FortiGate-6000 cluster that is operating normally. The output shows which FortiGate-6000 has become the primary FortiGate-6000 and how it was chosen. You can also see CPU and memory use data, HA heartbeat VLAN IDs, Chassis ID, and so on.

```
get system ha status
HA Health Status: OK
Model: FortiGate-6000F
Mode: HA A-P
Group Name: My-6K-cluster
Group ID : 6
Debug: 0
Cluster Uptime: 0 days 12:42:5
Cluster state change time: 2019-02-24 16:26:30
Primary selected using:
  <2019/02/24 16:26:30> F6KF31T018900143 is selected as the primary because it has the
largest value of serialno.
ses_pickup: disable
override: disable
Configuration Status:
  F6KF31T018900143(updated 3 seconds ago): in-sync
  F6KF31T018900143 chksum dump: 7c 74 ce 81 83 c0 54 c1 01 1d 4f a9 c9 fd 17 df
  F6KF51T018900022(updated 4 seconds ago): in-sync
  F6KF51T018900022 chksum dump: 7c 74 ce 81 83 c0 54 c1 01 1d 4f a9 c9 fd 17 df
System Usage stats:
  F6KF31T018900143(updated 4 seconds ago):
    sessions=198, average-cpu-user/nice/system/idle=1%/0%/0%/97%, memory=5%
  F6KF51T018900022 (updated 0 seconds ago):
    sessions=0, average-cpu-user/nice/system/idle=2%/0%/0%/96%, memory=6%
HBDEV stats:
  F6KF31T018900143(updated 4 seconds ago):
    ha1: physical/10000full, up, rx-bytes/packets/dropped/errors=227791977/902055/0/0,
tx=85589814/300318/0/0, vlan-id=4091
    ha2: physical/10000full, up, rx-bytes/packets/dropped/errors=227791977/902055/0/0,
tx=85589814/300318/0/0, vlan-id=4092
  F6KF51T018900022(updated 0 seconds ago):
    ha1: physical/10000full, up, rx-bytes/packets/dropped/errors=0/0/0/0,
tx=85067/331/0/0, vlan-id=4091
```

```

    ha2: physical/10000full, up, rx-bytes/packets/dropped/errors=947346/3022/0/0,
tx=206768/804/0/0, vlan-id=4092
Primary: 6K-Chassis-1    , F6KF31T018900143, cluster index = 0
Secondary: 6K-Chassis-2  , F6KF51T018900022, cluster index = 1
number of vcluster: 1
vcluster 1: work 10.101.11.20
Primary: F6KF31T018900143, operating cluster index = 0
Secondary: F6KF51T018900022, operating cluster index = 1
Chassis Status: (Local chassis ID: 2)
  Chassis ID 1: Secondary Chassis
    Slot ID 1: Primary Slot
    Slot ID 2: Secondary Slot
  Chassis ID 2: Primary Chassis
    Slot ID 1: Primary Slot
    Slot ID 2: Secondary Slot

```

Confirming that the FortiGate-6000 HA cluster is synchronized

After an HA cluster is up and running, you can use the HA Status dashboard widget to view status information about the cluster. You can also use the `get system ha status` command to confirm that the cluster is operating normally. As highlighted below, the command shows the HA health status, describes how the current primary FortiGate-6000 was selected, shows if the configuration is synchronized (configuration status), and indicates the serial numbers of the primary and secondary FortiGate-6000s.

```

get system ha status
HA Health Status: OK
...
Primary selected using:
 <2019/09/23 12:56:53> FG6KF43E17000073 is selected as the primary because it has the
largest value of override priority.
...

Configuration Status:
  FG6KF43E17000073(updated 2 seconds ago): in-sync
  FG6KF43E17000073 chksum dump: 0b 16 f2 e4 e2 89 eb a1 bf 8f 15 9b e1 4e 3b f2
  FG6KF43E17000065(updated 4 seconds ago): in-sync
  FG6KF43E17000065 chksum dump: 0b 16 f2 e4 e2 89 eb a1 bf 8f 15 9b e1 4e 3b f2
...
Primary  : FG6KF43E17000073, operating cluster index = 0
Secondary: FG6KF43E17000065, operating cluster index = 1

```

For a FortiGate-6000 HA cluster to operate normally, the configurations of both FortiGate-6000s and the management boards and all of the FPCs in these devices must be synchronized. The `Configuration Status` information provided by the `get system ha status` command is a useful indicator of synchronization status of the cluster. The information provided indicates whether the FortiGate-6000s in the cluster are `in-sync` (or `out-of-sync`) and includes checksums of each FortiGate-6000 configuration. If the two FortiGate-6000s are synchronized, these checksums must match.

Viewing more details about HA cluster synchronization

You can use the `diagnose sys ha checksum show` command to display the debugzone and configuration checksums for the FortiGate-6000 in the cluster that you have logged in to.

```
diagnose sys ha checksum show
is_manage_primary()=1, is_root_primary()=1
debugzone
global: 7e 06 79 02 65 a9 ea e3 68 58 73 c2 33 d0 16 f1
root: 43 2c ee 2c f1 b3 b2 13 ff 37 34 5e 86 11 dc bf
mgmt-vdom: 9c 7d 58 9f 81 4b b7 4e ed 2a c3 02 34 b4 7c 63
all: 0b 16 f2 e4 e2 89 eb a1 bf 8f 15 9b e1 4e 3b f2

checksum
global: 7e 06 79 02 65 a9 ea e3 68 58 73 c2 33 d0 16 f1
root: 43 2c ee 2c f1 b3 b2 13 ff 37 34 5e 86 11 dc bf
mgmt-vdom: 9c 7d 58 9f 81 4b b7 4e ed 2a c3 02 34 b4 7c 63
all: 0b 16 f2 e4 e2 89 eb a1 bf 8f 15 9b e1 4e 3b f2
```

The first line of this example output indicates that the command is displaying information for the primary FortiGate-6000. This command output then shows debugzone and checksum information for the management board. You can verify that the management board is synchronized because both sets of checksums match.

Each set of checksums includes a checksum for the global configuration, for each VDOM (in this case there are two VDOMs: root and mgmt-vdom), and a checksum for the complete configuration (all).

You can use the `diagnose sys ha checksum cluster` command to display the debugzone and configuration checksums for both FortiGate-6000s in the cluster. The command output also indicates which FortiGate-6000 is the primary (`is_manage_primary()=1`) and the secondary (`is_manage_primary()=0`). If the cluster is synchronized, both FortiGate-6000s will have the same checksums.

```
diagnose sys ha checksum cluster

===== F6KF31T018900158 =====

is_manage_primary()=0, is_root_primary()=0
debugzone
global: b7 df c2 39 be 5c 3f ac cb 6f 53 20 5a b6 2d 98
root: 1b 71 bc 50 80 15 10 5c 7e 79 38 73 30 dd 56 32
mgmt-vdom: 79 f5 78 e4 ad 6d 39 b8 8e 96 84 21 18 28 18 64
all: 49 63 81 37 c1 a2 78 95 46 44 08 ff 5d 2e 44 a7

checksum
global: b7 df c2 39 be 5c 3f ac cb 6f 53 20 5a b6 2d 98
root: 1b 71 bc 50 80 15 10 5c 7e 79 38 73 30 dd 56 32
mgmt-vdom: 79 f5 78 e4 ad 6d 39 b8 8e 96 84 21 18 28 18 64
all: 49 63 81 37 c1 a2 78 95 46 44 08 ff 5d 2e 44 a7

===== F6KF31T018900139 =====

is_manage_primary()=1, is_root_primary()=1
debugzone
global: b7 df c2 39 be 5c 3f ac cb 6f 53 20 5a b6 2d 98
root: 1b 71 bc 50 80 15 10 5c 7e 79 38 73 30 dd 56 32
mgmt-vdom: 79 f5 78 e4 ad 6d 39 b8 8e 96 84 21 18 28 18 64
all: 49 63 81 37 c1 a2 78 95 46 44 08 ff 5d 2e 44 a7
```

```
checksum
global: b7 df c2 39 be 5c 3f ac cb 6f 53 20 5a b6 2d 98
root: 1b 71 bc 50 80 15 10 5c 7e 79 38 73 30 dd 56 32
mgmt-vdom: 79 f5 78 e4 ad 6d 39 b8 8e 96 84 21 18 28 18 64
all: 49 63 81 37 c1 a2 78 95 46 44 08 ff 5d 2e 44 a7
```

Finally, you can also log into the CLI of each FortiGate-6000 in the cluster and use the `diagnose sys confsync showchecksum` command to confirm that the configurations of the management board and the FPCs in each FortiGate-6000 are synchronized.

The output of the command will also show that the ha checksums are the same for both FortiGate-6000s, but the confsync checksums are different. This occurs because some parts of the configuration are not synchronized by HA so each FortiGate-6000 will have a different configuration and different confsync checksums.

See [Viewing more details about FortiGate-6000 synchronization on page 21](#) for details about the `diagnose sys confsync showchecksum` command.

Primary FortiGate-6000 selection with override disabled (default)

FortiGate-6000 FGCP selects the primary FortiGate-6000 based on [standard FGCP primary unit selection](#) and also accounting for the number of failed FPCs. The selection sequence is:

- At least one active FPC
- Connected monitored interfaces
- Number of active FPCs
- Number of active SSDs (if SSD failure protection is enabled, FortiGate-6301F or 6501F only)
- Age
- Device priority
- Serial number

In most cases and with default settings, if everything is connected and operating normally, the FortiGate-6000 with the highest serial number becomes the primary FortiGate-6000. You can set the device priority higher on one of the FortiGate-6000s if you want it to become the primary FortiGate-6000.

The selection sequence also shows that at least one FPC must be active for a FortiGate-6000 to be selected to be the primary. If at least one FPC is active on each FortiGate-6000, the most important criteria is the number of connected monitored interfaces followed by the number of failed FPCs, followed by the number of active SSDs if SSD failure protection is enabled. So if one or more FPCs fail, if interface monitoring is not configured or no monitored interface has become disconnected, the primary FortiGate-6000 will be the one with the most active FPCs.

Primary FortiGate-6000 selection with override enabled

With override enabled, FortiGate-6000 FGCP selects the primary FortiGate-6000 based on [standard FGCP primary unit selection with override enabled](#) and also accounting for the number of failed FPCs. The selection sequence is:

- At least one active FPC
- Connected monitored interfaces
- Number of active FPCs

- Number of active SSDs (if SSD failure protection is enabled)
- Device priority
- Age
- Serial number

Enabling override and adjusting the device priority means that the FortiGate-6000 with the highest device priority becomes the primary FortiGate-6000 as long as monitored interfaces are not configured, no monitored interface has become disconnected, both FortiGate-6000s have the same number of failed FPCs, and SSD failure protection is either not enabled or an SSD has not failed. Enabling override causes the cluster to negotiate more often to make sure that the FortiGate-6000 with the highest device priority always becomes the primary FortiGate-6000.

Failover protection

FortiGate-6000 HA supports failover protection to provide FortiOS services even when one of the FortiGate-6000s encounters a problem that would result in partial or complete loss of connectivity or reduced performance for a standalone FortiGate-6000. This failover protection provides a backup mechanism that can be used to reduce the risk of unexpected downtime, especially in a mission-critical environment.

To achieve failover protection in a FortiGate-6000 cluster, one of the FortiGate-6000s functions as the primary, processing traffic and the other as the secondary, operating in an active stand-by mode. The cluster IP addresses and HA virtual MAC addresses are associated with the interfaces of the primary. All traffic directed at the cluster is actually sent to and processed by the primary.

While the cluster is functioning, the primary FortiGate-6000 functions as the FortiGate network security device for the networks that it is connected to. In addition, the primary FortiGate-6000 and the secondary FortiGate-6000 use the HA heartbeat to keep in constant communication. The secondary FortiGate-6000 reports its status to the primary FortiGate-6000 and receives and stores connection and state table updates from the primary FortiGate-6000.

FortiGate-6000 HA supports three kinds of failover protection:

- Device failure protection automatically replaces a failed device and restarts traffic flow with minimal impact on the network.
- Link failure protection maintains traffic flow if a link fails.
- FPC failure protection makes sure that traffic is processed by the FortiGate-6000 with the most operating FPCs.
- Session failure protection resumes communication sessions with minimal loss of data if a device, FPC, or link failure occurs.

Device failure

If the primary FortiGate-6000 encounters a problem that is severe enough to cause it to fail, the secondary FortiGate-6000 becomes new primary FortiGate-6000. This occurs because the secondary FortiGate-6000 is constantly waiting to negotiate to become primary FortiGate-6000. Only the heartbeat packets sent by the primary FortiGate-6000 keep the secondary FortiGate-6000 from becoming the primary FortiGate-6000. Each received heartbeat packet resets a negotiation timer in the secondary FortiGate-6000. If this timer is allowed to run out because the secondary FortiGate-6000 does not receive heartbeat packets from the primary FortiGate-6000, the secondary FortiGate-6000 assumes that the primary FortiGate-6000 has failed and becomes the primary FortiGate-6000.

The new primary FortiGate-6000 will have the same MAC and IP addresses as the former primary FortiGate-6000. The new primary FortiGate-6000 then sends gratuitous ARP packets out all of its connected interfaces to inform attached switches to send traffic to the new primary FortiGate-6000. Sessions then resume with the new primary FortiGate-6000.

Link failure

If your HA configuration includes HA interface monitoring, if a primary FortiGate-6000 interface fails or is disconnected while a cluster is operating, a link failure occurs. When a link failure occurs, the FortiGate-6000s in the cluster negotiate to select a new primary FortiGate-6000. The link failure means that a that primary FortiGate-6000 with the most link failures will become the secondary and the FortiGate-6000 with the fewest link failures becomes the primary FortiGate-6000.

Just as for a device failover, the new primary FortiGate-6000 sends gratuitous arp packets out all of its connected interfaces to inform attached switches to send traffic to it. Sessions then resume with the new primary FortiGate-6000.

If the secondary FortiGate-6000 experiences a link failure, its status in the cluster does not change. However, in future negotiations FortiGate-6000 with a link failure is less likely to become the primary FortiGate-6000.

If one of the FortiGate-6000s experiences an FPC failure and the other experiences a link failure, the FortiGate-6000 with the most operating FPCs becomes the primary FortiGate-6000, even if it is also experiencing a link failure.

FPC failure

If one or more FPCs in the primary FortiGate-6000 fails, the cluster renegotiates and the FortiGate-6000 with the most operating FPCs becomes the primary FortiGate-6000. An FPC failure can occur if an FPC shuts down due to a software crash or hardware problem, or if the FPC is manually shut down.

FPCs also shut down if two of the three FortiGate-6000 power supply units (PSUs) become disconnected from their power source. The FortiGate-6000 includes three hot-swappable PSUs in a 2+1 redundant configuration. At least two of the PSUs must be operating to provide power to the FortiGate-6000. If only one PSU is operating, only four of the FPCs will continue running (usually the FPCs in slots 1 to 4). For more information about FPC failure with power loss, see [AC power supply units \(PSUs\)](#).

From the management board GUI dashboard, the Sensor Information dashboard widget displays information about the status of the power supplies. If all power supplies are operating, the widget displays their **Status** as **Normal**.

From the management board CLI, you can use the `execute sensor list` command to verify if the power supplies are operating. The command displays the current status of all FortiGate-6000 sensors including the power supply sensors. Power supply sensor entries should be similar to the following (shown for a FortiGate-6301E). The power supply sensor lines start with `PS{1|2|3}`:

```
65 PS1 VIN          alarm=0 value=122 threshold_status=0
66 PS1 VOUT_12V     alarm=0 value=12.032 threshold_status=0
67 PS1 Temp 1       alarm=0 value=24 threshold_status=0
68 PS1 Temp 2       alarm=0 value=36 threshold_status=0
69 PS1 Fan 1        alarm=0 value=8832 threshold_status=0
70 PS1 Status       alarm=0
71 PS2 VIN          alarm=0 value=122 threshold_status=0
72 PS2 VOUT_12V     alarm=0 value=12.032 threshold_status=0
73 PS2 Temp 1       alarm=0 value=24 threshold_status=0
74 PS2 Temp 2       alarm=0 value=37 threshold_status=0
75 PS2 Fan 1        alarm=0 value=9088 threshold_status=0
76 PS2 Status       alarm=0
```

```

77 PS3 VIN          alarm=0 value=122 threshold_status=0
78 PS3 VOUT_12V     alarm=0 value=12.032 threshold_status=0
79 PS3 Temp 1       alarm=0 value=23 threshold_status=0
80 PS3 Temp 2       alarm=0 value=37 threshold_status=0
81 PS3 Fan 1        alarm=0 value=9088 threshold_status=0
82 PS3 Status       alarm=0

```

Any non zero alarm or threshold_status values indicate a possible problem with that power supply.

After the primary FortiGate-6000 in an HA cluster experiences an FPC failure, the cluster negotiates and the FortiGate-6000 with the most operating FPCs becomes the new primary FortiGate-6000. The new primary FortiGate-6000 sends gratuitous arp packets out all of its connected interfaces to inform attached switches to send traffic to it. Sessions then resume with the new primary FortiGate-6000.

If the secondary FortiGate-6000 experiences an FPC failure, its status in the cluster does not change. In future cluster negotiations the FortiGate-6000 with an FPC failure is less likely to become the primary FortiGate-6000.



To prevent multiple failovers, if an FPC failure occurs in an HA cluster with override enabled, you should disable override until you can fix the problems and get all the FPCs up and running and synchronized.

After an FPC failure, sessions and configuration changes are not synchronized to the failed FPCs.

If failed FPCs recover in the secondary FortiGate-6000, it will continue to operate as the secondary FortiGate-6000 and will attempt to re-synchronize the FPCs with the management board. This process may take a few minutes, but if it is successful, the secondary FortiGate-6000 can return to fully participate in the cluster.

If there have been many configuration changes, the FPCs need to be manually synchronized with the management board. Log into the CLI of each out of synch FPC and enter the `execute factoryreset` command to reset the configuration. After the FPC restarts, the management board will attempt to synchronize its configuration. If the configuration synchronization is successful, the FPC can start processing traffic again.

If there has been a firmware upgrade, and the firmware running on the failed FPC is out of date, you can upgrade the firmware of the FPC as described in the section: [Installing firmware on an individual FPC on page 102](#).

You can optionally use the following command to make sure the sessions on the FPCs in the secondary FortiGate-6000 are synchronized with the sessions on the FPCs in the primary FortiGate-6000.

```
diagnose test application chlbd 10
```

Once all of the FPCs are operating and synchronized, the secondary FortiGate-6000 can fully participate with the cluster.

For more information about troubleshooting FPC failures, see [Troubleshooting an FPC failure on page 108](#).

SSD failure

FortiGate-6501F or 6301F HA clusters support SSD (log disk) failure protection. SSD failure protection is disabled by default. You can use the following command to enable SSD failure protection:

```

config system ha
    set ssd-failover enable
end

```

If SSD failure detection is enabled, if an SSD on the primary FortiGate-6501F or 6301F fails, an HA failover occurs and the FortiGate-6501F or 6301F with the most operating SSDs becomes the primary.

If an SSD fails on the secondary FortiGate-6501F or 6301F, its status in the cluster does not change. However, in future negotiations the FortiGate-6501F or 6301F with an the most SSD failures is less likely to become the primary.

Session failover

If you enable session failover (also called session pickup) for the cluster, during cluster operation the primary FortiGate-6000 informs the secondary FortiGate-6000 of changes to the primary FortiGate-6000 connection and state tables, keeping the secondary FortiGate-6000 up-to-date with the traffic currently being processed by the cluster.

Session synchronization traffic uses the HA1 and HA2 interfaces. FortiGate-6000 does not support using the `session-sync-dev` option to use data interfaces for session synchronization. The HA1 and HA2 interfaces provide enough bandwidth for both HA heartbeat and session synchronization traffic, so additional session synchronization devices are not required. As well, keeping session synchronization traffic on the HA1 and HA2 interfaces separates session synchronization traffic from data traffic.

After a failover the new primary FortiGate-6000 recognizes open sessions that were being handled by the cluster. The sessions continue to be processed by the new primary FortiGate-6000 and are handled according to their last known state.

If you leave session pickup disabled, the cluster does not keep track of sessions and after a failover, active sessions have to be restarted or resumed.

Primary FortiGate-6000 recovery

If a primary FortiGate-6000 recovers after a device, FPC, or link failure, it will operate as a subordinate unit. If `override` is enabled; however, when the FortiGate-6000 recovers, the cluster will renegotiate and the FortiGate-6000 with the highest device priority becomes the primary.

HA reserved management interfaces

You can edit an HA cluster and configure one or more of the interfaces in the `mgmt-vdom` VDOM (`mgmt1`, `mgmt2`, and `mgmt3`) to be HA reserved management interfaces. You can then log into each FortiGate-6000 in the cluster and configure its reserved management interfaces with IP addresses and other custom interface settings as required. You can also configure routing for each reserved management interface. The result is that each FortiGate-6000 in the cluster has its own management interface or interfaces and each of these interfaces has its own IP address that is not synchronized to the other FortiGate-6000 in the cluster.

To configure an HA reserved management interface from the GUI, go to **System > HA** and enable **Management Interface Reservation**. Select one or more interfaces to be HA reserved management interfaces. Optionally configure routing for each reserved management interface. This routing configuration is not synchronized and can be configured separately for each FortiGate-6000 in the cluster.

To configure an HA reserved management interface from the CLI:

```
config system ha
  set mode a-p
  set ha-mgmt-status enable
  set ha-direct enable
  config ha-mgmt-interfaces
    edit 0
```

```

    set interface <interface>
    set dst <destination-ip>
    set gateway <gateway-ip>
    set gateway6 <gateway-ipv6-ip>
end
end

```

Enabling `ha-direct` from the CLI is required if you plan to use the HA reserved management interface for SNMP, remote logging, or communicating with FortiSandbox. Enabling `ha-direct` is also required for some types of remote authentication, but is not required for RADIUS remote authentication.

<interface> can be `mgmt1`, `mgmt2`, or `mgmt3`. You can only select an interface if it has not been used in another configuration.

For more information, see [Out-of-band management](#).

HA in-band management for management interfaces

The FortiGate 6000F supports [FGCP HA in-band management](#) for FortiGate 6000F management interfaces (`mgmt1`, `mgmt2`, and `mgmt3`).

HA in-band management allows you to add a second management IP address to `mgmt1`, `mgmt2`, or `mgmt3`. The management IP address is accessible from the network that the interface is connected to. This setting is not synchronized, so both FortiGate 6000Fs in the cluster can have their own in-band management IP addresses; providing management access to the secondary FortiGate 6000F.



FortiGate 6000F does not support HA in-band management for data interfaces.

FortiGate-6000 HA in-band management configuration:

```

config vdom
  edit mgmt-vdom
    config system interface
      edit {1-mgmt1 | 1-mgmt2 | 1-mgmt3 | 2-mgmt1 | 2-mgmt2 | 2-mgmt3}
        set management-ip <ip address>
      end
    end
  end
end

```

The `management-ip` option is available only when HA is enabled.

To support HA in-band management, the FortiGate 6000F handles [HA virtual MAC addresses](#) in the same way as other FortiGates.

Virtual clustering

A virtual cluster consists of two FortiGate 6000Fs operating in active-passive HA mode with Multi VDOM mode enabled. Virtual clustering is an extension of FGCP HA that uses VDOM partitioning to send traffic for some VDOMs to the primary FortiGate 6000F and traffic for other VDOMs to the secondary FortiGate 6000F. Distributing traffic between the

FortiGate 6000Fs in a virtual cluster is similar to load balancing and can potentially improve overall throughput. You can adjust VDOM partitioning at any time to optimize traffic distribution without interrupting traffic flow.

VDOM partitioning distributes VDOMs between two virtual clusters (virtual cluster 1 and virtual cluster 2). When configuring virtual clustering you would normally set the device priority of virtual cluster 1 higher for the primary FortiGate 6000F and the device priority of virtual cluster 2 higher for the secondary FortiGate 6000F. With this configuration, all traffic in the VDOMs in virtual cluster 1 is processed by the primary FortiGate 6000F and all traffic in the VDOMs in virtual cluster 2 is processed by the secondary FortiGate 6000F. The FGCP selects the primary and secondary FortiGate 6000F whenever the cluster negotiates. The primary FortiGate 6000F can dynamically change based on FGCP HA primary unit selection criteria.

If a failure occurs and only one FortiGate 6000F continues to operate, all traffic fails over to that FortiGate 6000F, similar to normal FGCP HA. When the failed FortiGate 6000F rejoins the cluster, the configured traffic distribution is restored.

For more information about virtual clustering see [HA virtual cluster setup](#).



If you don't want active-passive virtual clustering to distribute traffic between FortiGate 6000Fs, you can configure VDOM partitioning to send traffic for all VDOMs to the primary FortiGate 6000F. The result is the same as standard active-passive FGCP HA, all traffic is processed by the primary FortiGate 6000F.

Virtual clustering creates a cluster between instances of each VDOM on the two FortiGate 6000Fs in the virtual cluster. All traffic to and from a given VDOM is sent to one of the FortiGate 6000Fs where it stays within its VDOM and is only processed by that VDOM. One FortiGate 6000F is the primary FortiGate 6000F for each VDOM and one FortiGate 6000F is the secondary FortiGate 6000F for each VDOM. The primary FortiGate 6000F processes all traffic for its VDOMs. The secondary FortiGate 6000F processes all traffic for its VDOMs.

The HA heartbeat and session synchronization provides the same HA services in a virtual clustering configuration as in a standard HA configuration. One set of HA heartbeat interfaces provides HA heartbeat and session synchronization services for all of the VDOMs in the cluster. You do not have to add a heartbeat interface for each VDOM.

Limitations of FortiGate-6000 virtual clustering

FortiGate-6000 virtual clustering includes the following limitations:

- Virtual clustering supports two FortiGate 6000Fs only.
- Active-passive HA mode is supported, active-active HA is not.
- The root and mgmt-vdom VDOMs must be in virtual cluster 1 (also called the primary virtual cluster).
- A VLAN must be in the same virtual cluster as the physical interface or LAG that the VLAN has been added to. The VLAN can be in the same VDOM as its physical interface or LAG or in a different VDOM, as long as both VDOMs are in the same virtual cluster.
- The interfaces that are created when you add an inter-VDOM link must be in the same virtual cluster as the inter-VDOM link. You can change the virtual cluster that an inter-VDOM link is in by editing the inter-VDOM link and changing the `vcluster` setting.

Virtual clustering VLAN/VDOM limitation

In a FortiGate 6000F virtual clustering configuration, a VLAN must be in the same virtual cluster as the physical interface, LAG, or redundant interface that the VLAN has been added to. The VLAN can be in the same VDOM as its physical interface, LAG, or redundant interface or in a different VDOM, as long as both VDOMs are in the same virtual cluster.

If virtual clustering has already been set up, when adding VLANs, GUI and CLI error checking prevents you from adding a VLAN to a VDOM that is in a different virtual cluster than the physical interface, LAG, or redundant interface that you are attempting to add the VLAN to. However, error checking can't prevent this problem if you configure the VLANs before setting up virtual clustering or if you move VDOMs to different virtual clusters after adding the VLANs.

A recommended strategy for preventing this problem could involve the following steps:

1. Start by setting up virtual clustering before creating new VDOMs.
2. Create a placeholder VDOM and add it to virtual cluster 2.
3. Separate traffic interfaces between the root VDOM in virtual cluster 1 and the placeholder VDOM in virtual cluster 2. Based on network planning you can create an even distribution of planned traffic volume between the two virtual clusters.
4. Build up your configuration by adding more VDOMs, LAGs, redundant interfaces, and VLANs as required, making sure to keep VLANs in the same virtual cluster as their parent interfaces, LAGs, or redundant interfaces.

Example incorrect VLAN configuration

Consider the following FortiGate 6000F virtual clustering example, which shows how traffic can be blocked by this limitation:

- Three data traffic VDOMs: root, Engineering, and Marketing.
- One LAG interface: LAG1 in the root VDOM.
- Two VLAN interfaces added to LAG1: vlan11 and vlan12.
 - vlan11 is added to the Engineering VDOM.
 - vlan12 is added to the Marketing VDOM.
- The root and Engineering VDOMs are in virtual cluster 1.
- The Marketing VDOM is in virtual cluster 2.

As a result of this configuration:

- vlan11 is in the Engineering VDOM, which is in virtual cluster 1. vlan11 is also in LAG1, which is in the root VDOM, also in virtual cluster 1. vlan11 and its LAG are in the same virtual cluster. Traffic can pass through vlan11.
- vlan12 is in the Marketing VDOM, which is in virtual cluster 2. vlan12 is also in LAG1, which is in the root VDOM, in virtual cluster 1. vlan12 and its LAG are in different virtual clusters. Traffic cannot pass through vlan12.

Configuring virtual clustering

Configuring virtual clustering is the same as configuring standard FCGP HA with the addition of VDOM partitioning. Using VDOM partitioning, you can control the distribution of VDOMs, and the traffic they process, between the FortiGates in the cluster.

VDOM partitioning can be thought of in two parts. First, there is configuring the distribution of VDOMs between two virtual clusters. By default, all VDOMS are in virtual cluster 1, virtual cluster 1 is associated with the primary FortiGate 6000F, and the primary FortiGate 6000F processes all traffic. If you want traffic to be processed by the secondary FortiGate 6000F, you need to enable virtual cluster 2, move some of the VDOMs to it, and associate virtual cluster 2 with the secondary FortiGate 6000F.

You associate a virtual cluster with a FortiGate 6000F using device priorities. The FortiGate 6000F with the highest device priority is associated with virtual cluster 1. To associate a FortiGate 6000F with virtual cluster 2, you must enable virtual cluster 2 and set virtual cluster 2 device priorities on each FortiGate 6000F. The FortiGate 6000F with the highest

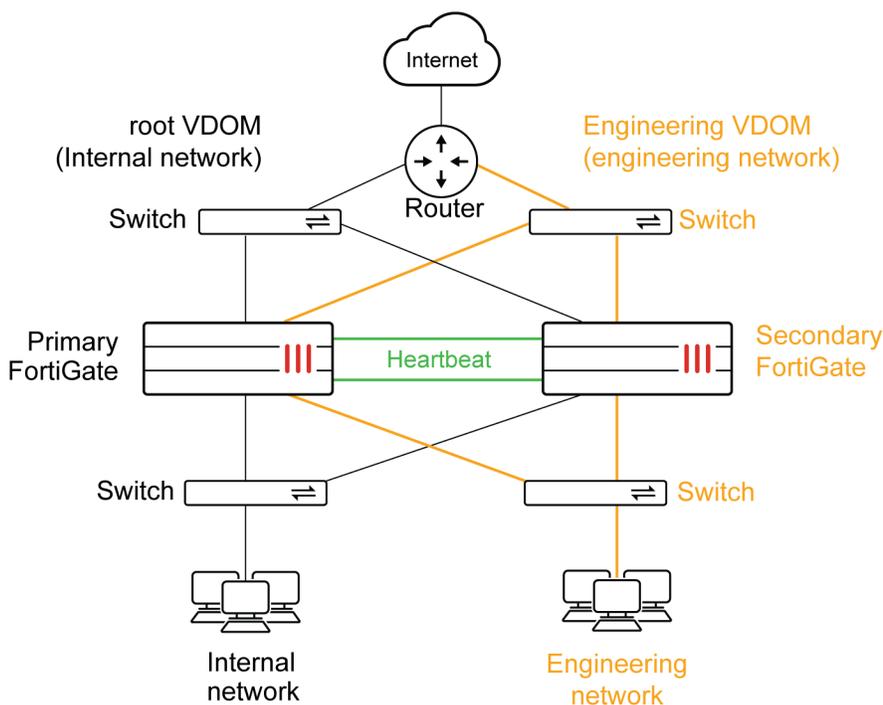
virtual cluster 2 device priority processes traffic for the VDOMs added to virtual cluster 2. (Reminder: device priorities are not synchronized.)

Normally, you would set the virtual cluster 1 device priority for the primary FortiGate 6000F and the virtual cluster 2 device priority higher for the secondary FortiGate 6000F. Then the primary FortiGate 6000F would process virtual cluster 1 traffic and the secondary FortiGate 6000F would process virtual cluster 2 traffic.

Enabling virtual cluster 2 also turns on HA override for virtual cluster 1 and 2. Enabling override is required for virtual clustering to function as configured. Enabling override causes the cluster to negotiate every time the cluster state changes. If override is not enabled, the cluster may not negotiate as often. While more frequent negotiation may cause more minor traffic disruptions, with virtual clustering its more important to negotiate after any state change to make sure the configured traffic flows are maintained.

The figure below shows a simple FortiGate 6000F virtual cluster that provides redundancy and failover for two networks. The configuration includes two VDOMs. The root VDOM handles internal network traffic and the Engineering VDOM handles Engineering network traffic. VDOM partitioning has been set up to send all root VDOM traffic to the primary FortiGate and all Engineering VDOM traffic to the secondary FortiGate.

Example virtual clustering configuration



Primary FortiGate 6000F configuration

The primary FortiGate 6000F configuration:

- Sets the primary FortiGate 6000F to be chassis 1.
- Enables virtual cluster 2 (`vcluster2`) to enable virtual clustering.
- Enables override for virtual cluster 1.
- Sets the virtual cluster 1 device priority to 200.
- Enables override for virtual cluster 2 (`secondary-vcluster`).

- Sets the virtual cluster 2 device priority to 50.
- Adds the Engineering VDOM to virtual cluster 2 (all VDOMs remain in virtual cluster 1 unless you add them to virtual cluster 2).

```
config system ha
  set group-id 6
  set group-name <name>
  set mode a-p
  set password <password>
  set hbdev ha1 50 ha2 50
  set chassis-id 1
  set vcluster2 enable
  set override enable
  set priority 200
config secondary-vcluster
  set override enable
  set priority 50
  set vdom Engineering
end
```

Secondary FortiGate configuration

The secondary FortiGate configuration:

- Sets the secondary FortiGate to be chassis 2.
- Enables virtual cluster 2 (`vcluster2`) to enable virtual clustering.
- Enables override for virtual cluster 1.
- Sets the device priority of virtual cluster 1 to 50.
- Enables override for virtual cluster 2 (`secondary-vcluster`).
- Sets the virtual cluster 2 device priority to 200.
- You do not need to add the Engineering VDOM to virtual cluster 2, the configuration of the VDOMs in virtual cluster 2 is synchronized from the primary FortiGate.

```
config system ha
  set group-id 6
  set group-name <name>
  set mode a-p
  set password <password>
  set hbdev "ha1" 50 "ha2" 50
  set chassis-id 2
  set vcluster2 enable
  set override enable
  set priority 50
config secondary-vcluster
  set override enable
  set priority 200
  set vdom Engineering
end
```



Since the primary FortiGate 6000F has the highest device priority, it processes all traffic for the VDOMs in virtual cluster 1. Since the secondary FortiGate 6000F has the highest virtual cluster 2 device priority, it processes all traffic for the VDOM in virtual cluster 2. The primary FortiGate 6000F configuration adds the VDOMs to virtual cluster 2. All you have to configure on the secondary FortiGate 6000F for virtual cluster 2 is the virtual cluster 2 (or secondary-vcluster) device priority.

Virtual cluster GUI configuration

From the GUI, you configure virtual clustering from the **Global** menu by going to **System > HA**, configuring HA settings and VDOM Partitioning.

Primary FortiGate VDOM partitioning

VDOM Partitioning

Virtual cluster 1	<div style="border: 1px solid #ccc; padding: 2px;"> mgmt-vdom root ✕ </div>
	+
Virtual cluster 2	<div style="border: 1px solid #ccc; padding: 2px;"> Engineering ✕ </div>
	+

Secondary Cluster Settings

Device priority i

Secondary FortiGate VDOM partitioning

VDOM Partitioning

Virtual cluster 1	<div style="border: 1px solid #ccc; padding: 2px;"> mgmt-vdom root ✕ </div>
	+
Virtual cluster 2	<div style="border: 1px solid #ccc; padding: 2px;"> Engineering ✕ </div>
	+

Secondary Cluster Settings

Device priority i

HA cluster firmware upgrades

Both management boards and all of the FPCs in a FortiGate 6000F HA cluster run the same firmware image. You upgrade the firmware from the primary FortiGate 6000F management board.

You can perform a graceful firmware upgrade of an FGCP cluster by setting `upgrade-mode` to `uninterruptible` and enabling `session-pickup`. A graceful firmware upgrade only causes minimal traffic interruption. Use the following command to enable these settings; they are disabled by default. These settings are synchronized.

```
config system ha
  set upgrade-mode uninterruptible
  set session-pickup enable
end
```

When these settings are enabled, the primary FortiGate 6000F management board uploads firmware to the secondary FortiGate 6000F management board. The secondary management board uploads the firmware to all of the FPCs in the secondary FortiGate 6000F. Then the management board and all of the FPCs in the secondary FortiGate 6000F upgrade their firmware, reboot, and resynchronize.

Then all traffic fails over to the secondary FortiGate 6000F which becomes the new primary FortiGate 6000F. Then the management board and the FPCs in the new secondary FortiGate 6000F upgrade their firmware and rejoin the cluster. Unless `override` is enabled, the new primary FortiGate 6000F continues to operate as the primary FortiGate 6000F.

Normally you would want to set `upgrade-mode` to `uninterruptible` to minimize traffic interruptions. But `upgrade-mode` does not have to be set to `uninterruptible`. In fact, if a traffic interruption is not going to cause any problems, you can set `upgrade-mode` to `simultaneous` so that the firmware upgrade process takes less time.

As well, some firmware upgrades may not support setting `upgrade-mode` to `uninterruptible`. Make sure to review the release notes before running a firmware upgrade to verify whether or not setting `upgrade-mode` to `uninterruptible` is supported to upgrade to that version.



To make sure a FortiGate 6000F firmware upgrade is successful, before starting the upgrade Fortinet recommends you use health checking to make sure the management board and the FPCs are all synchronized and operating as expected.

If you are following a multi-step upgrade path, you should re-do health checking after each upgrade step to make sure all components are synchronized before the next step.

You should also perform a final round of health checking after the firmware upgrade process is complete.

For recommended health checking commands, see the following Fortinet community article:

[Technical Tip: FortiGate-6000/7000 Chassis health check commands.](#)

Distributed clustering

FortiGate-6000 HA supports separating the FortiGate-6000s in an HA cluster to different physical locations. Distributed FortiGate-6000 HA clustering (or geographically distributed FortiGate-6000 HA or geo clustering) can involve two FortiGate-6000s in different rooms in the same building, different buildings in the same location, or even different geographical sites such as different cities, countries, or continents.

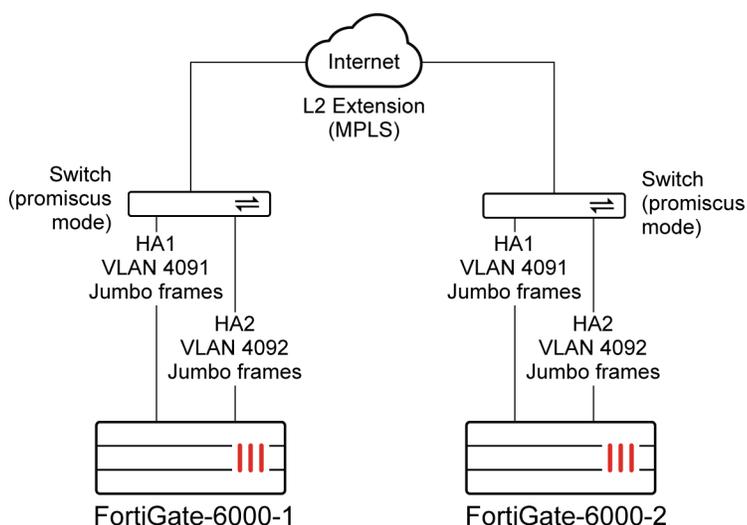
Just like any FortiGate-6000 HA configuration, distributed FortiGate-6000 HA requires heartbeat communication between the FortiGate-6000s over the HA1 and HA2 interfaces. In a distributed FortiGate-6000 HA configuration this heartbeat communication can take place over the internet or over other transmission methods including satellite linkups.

Most Data Center Interconnect (DCI) or MPLS-based solutions that support layer 2 extensions and VLAN tags between the remote data centers should also support HA heartbeat communication between the FortiGates in the distributed locations. Using VLANs and switches in promiscuous mode to pass all traffic between the locations can also be helpful.

You cannot change HA heartbeat IP addresses, so the heartbeat interfaces have to be able to communicate over the same subnet.

The HA1 and HA2 interface traffic must be separated. You can do this by using separate channels for each interface or by configuring the HA1 and HA2 interfaces to use different VLANs.

Example FortiGate-6000 distributed clustering configuration



Because of the possible distance between sites, it may take a relatively long time for heartbeat packets to be transmitted between the FortiGate-6000s. This could lead to a split brain scenario. To avoid a split brain scenario you can modify heartbeat timing so that the cluster expects extra time between heartbeat packets. As a general rule, set the heartbeat failover time (`hb-interval`) to be longer than the max latency or round trip time (RTT). You could also increase the `hb-lost-threshold` to tolerate losing heartbeat packets if the network connection is less reliable.

In addition you could use different link paths for heartbeat packets to optimize HA heartbeat communication. You could also configure QoS on the links used for HA heartbeat traffic to make sure heartbeat communication has the highest priority.

For information about changing the heartbeat interval and other heartbeat timing related settings, see [Modifying heartbeat timing on page 75](#).

Modifying heartbeat timing

If the FortiGate-6000s in the HA cluster do not receive heartbeat packets on time, the FortiGate-6000s in the HA configuration may each determine that the other FortiGate-6000 has failed. HA heartbeat packets may not be sent on time because of network issues. For example, if the HA1 and HA2 communications links between the FortiGate-6000s

become too busy to handle the heartbeat traffic. Also, in a distributed clustering configuration the round trip time (RTT) between the FortiGate-6000s may be longer the expected time between heartbeat packets.

In addition, if the FortiGate-6000s becomes excessively busy, they may delay sending heartbeat packets.

Even with these delays, the FortiGate-6000 HA cluster can continue to function normally as long as the HA heartbeat configuration supports longer delays between heartbeat packets and more missed heartbeat packets.

You can use the following commands to configure heartbeat timing:

```
config system ha
  set hb-interval <interval_integer>
  set hb-lost-threshold <threshold_integer>
  set hello-holddown <holddown_integer>
end
```

Changing the heartbeat interval

The heartbeat interval is the time between sending HA heartbeat packets. The heartbeat interval range is 1 to 20 (100*ms). The heartbeat interval default is 2 (200 ms).

A heartbeat interval of 2 means the time between heartbeat packets is 200 ms. Changing the heartbeat interval to 5 changes the time between heartbeat packets to 500 ms (5 * 100ms = 500ms).

Use the following CLI command to increase the heartbeat interval to 10:

```
config system ha
  set hb-interval 10
end
```

Changing the lost heartbeat threshold

The lost heartbeat threshold is the number of consecutive heartbeat packets that a FortiGate-6000 does not receive before assuming that a failure has occurred. The default value of 6 means that if a FortiGate-6000 does not receive 6 heartbeat packets, it determines that the other FortiGate-6000 in the cluster has failed. The range is 1 to 60 packets.

The lower the `hb-lost-threshold`, the faster a FortiGate-6000 HA configuration responds when a failure occurs. However, sometimes heartbeat packets may not be received because the other FortiGate-6000 is very busy or because of network conditions. This can lead to a false positive failure detection. To reduce these false positives you can increase the `hb-lost-threshold`.

Use the following command to increase the lost heartbeat threshold to 12:

```
config system ha
  set hb-lost-threshold 12
end
```

Adjusting the heartbeat interval and lost heartbeat threshold

The heartbeat interval combines with the lost heartbeat threshold to set how long a FortiGate-6000 waits before assuming that the other FortiGate-6000 has failed and is no longer sending heartbeat packets. By default, if a FortiGate-6000 does not receive a heartbeat packet from a cluster unit for $6 * 200 = 1200$ milliseconds or 1.2 seconds the FortiGate-6000 assumes that the other FortiGate-6000 has failed.

You can increase both the heartbeat interval and the lost heartbeat threshold to reduce false positives. For example, increasing the heartbeat interval to 20 and the lost heartbeat threshold to 30 means a failure will be assumed if no heartbeat packets are received after $30 * 2000$ milliseconds = 60,000 milliseconds, or 60 seconds.

Use the following command to increase the heartbeat interval to 20 and the lost heartbeat threshold to 30:

```
config system ha
  set hb-lost-threshold 20
  set hb-interval 30
end
```

Changing the time to wait in the hello state

The hello state hold-down time is the number of seconds that a FortiGate-6000 waits before changing from the hello state to the work state. After a failure or when starting up, FortiGate-6000s in HA mode operate in the hello state to send and receive heartbeat packets, to find each other, and form a cluster. A FortiGate-6000 should change from the hello state to the work state after it finds the FortiGate-6000 to form a cluster with. If for some reason the FortiGate-6000s cannot find each other during the hello state, both FortiGate-6000s may assume that the other one has failed and each could form separate clusters of one FortiGate-6000. The FortiGate-6000s could eventually find each other and negotiate to form a cluster, possibly causing a network interruption as they re-negotiate.

One reason for a delay of the FortiGate-6000s finding each other could be the FortiGate-6000s are located at different sites or for some other reason communication is delayed between the heartbeat interfaces. If you find that your FortiGate-6000s leave the hello state before finding each other you can increase the time that they wait in the hello state. The hello state hold-down time range is 5 to 300 seconds. The hello state hold-down time default is 20 seconds.

Use the following command to increase the time to wait in the hello state to 1 minute (60 seconds):

```
config system ha
  set hello-holddown 60
end
```

Changing how long routes stay in a cluster unit routing table

You can use the HA route time to live (`route-ttl`) option to control how long routes remain active in the new primary FortiGate-6000 after an FGCP HA failover. The default `route-ttl` is 600 seconds. The range is 5 to 3600 seconds (one hour). You can use the following command to change the `route-ttl` time.

```
config system ha
  set route-ttl <time>
end
```



FortiOS 6.0.6 for FortiGate-6000 does not support the `route-wait` and `route-hold` options.

To maintain communication sessions through a new primary FortiGate-6000, routes remain active in the routing table for the `route-ttl` time while the new primary FortiGate-6000 acquires new routes. Normally keeping `route-ttl` to the default value of 600 seconds (10 minutes) is acceptable because acquiring new routes and populating the routing tables of multiple FPCs can take a few minutes.

If the primary FortiGate-6000 needs to acquire a very large number of routes, or if for other reasons there is a delay in acquiring all routes, the primary FortiGate-6000 may not be able to maintain all communication sessions after a failover.

You can increase the `route-ttl` time if you find that communication sessions are lost after a failover. Increasing the `route-ttl` time allows the primary unit to use synchronized routes that are already in the routing table for a longer period of time while waiting to acquire new routes.

For more information, see [Synchronizing kernel routing tables](#).

Session failover (session-pickup)

Session failover means that after a failover, communication sessions resume on the new primary FortiGate-6000 with minimal or no interruption. Two categories of sessions need to be resumed after a failover:

- Sessions passing through the cluster
- Sessions terminated by the cluster

If sessions pickup is enabled, during cluster operation the primary FortiGate-6000 informs the secondary FortiGate-6000 of changes to the primary FortiGate-6000 connection and state tables for TCP and UDP sessions passing through the cluster, keeping the secondary FortiGate-6000 up-to-date with the traffic currently being processed by the cluster.

After a failover, the new primary FortiGate-6000 recognizes open sessions that were being handled by the cluster. The sessions continue to be processed by the new primary FortiGate-6000 and are handled according to their last known state.



Session-pickup has some limitations. For example, the FGCP does not support session failover for sessions being scanned by proxy-based security profiles. Session failover is supported for sessions being scanned by flow-based security profiles; however, flow-based sessions that fail over are not inspected after they fail over.

Sessions terminated by the cluster include management sessions (such as HTTPS connections to the FortiGate GUI or SSH connection to the CLI as well as SNMP and logging, and so on). Also included in this category are IPsec and SSL VPN sessions terminated by the cluster and explicit proxy sessions. In general, whether or not session-pickup is enabled, these sessions do not failover and have to be restarted.

Enabling session pickup for TCP SCTP and connectionless sessions

To enable session synchronization for TCP and SCTP sessions, enter:

```
config system ha
  set session-pickup enable
end
```

Turning on session synchronization for TCP and SCTP sessions by enabling `session-pickup` also turns on session synchronization for connectionless sessions, such as ICMP and UDP, by enabling `session-pickup-connectionless`. You can now choose to reduce processing overhead by not synchronizing connectionless sessions if you don't need to. If you want to synchronize connectionless sessions you can enable `session-pickup-connectionless`.

When `session-pickup` is enabled, sessions in the primary FortiGate-6000 TCP and connectionless session tables are synchronized to the secondary FortiGate-6000. As soon as a new session is added to the primary FortiGate-6000

session table, that session is synchronized to the secondary FortiGate-6000. This synchronization happens as quickly as possible to keep the session tables synchronized.

If the primary FortiGate-6000 fails, the new primary FortiGate-6000 uses its synchronized session tables to resume all TCP and connectionless sessions that were being processed by the former primary FortiGate-6000 with only minimal interruption. Under ideal conditions, all sessions should be resumed. This is not guaranteed though and under less than ideal conditions some sessions may need to be restarted.

If session pickup is disabled

If you disable session pickup, the FortiGate-6000 HA cluster does not keep track of sessions and after a failover, active sessions have to be restarted or resumed. Most session can be resumed as a normal result of how TCP and UDP resumes communication after any routine network interruption.



The session-pickup setting does not affect session failover for sessions terminated by the cluster.

If you do not require session failover protection, leaving session pickup disabled may reduce CPU usage and reduce HA heartbeat network bandwidth usage. Also, if your FortiGate-6000 HA cluster is mainly being used for traffic that is not synchronized (for example, for proxy-based security profile processing) enabling session pickup is not recommended since most sessions will not be failed over anyway.

Reducing the number of sessions that are synchronized

If session pickup is enabled, as soon as new sessions are added to the primary unit session table they are synchronized to the other cluster units. Enable the session-pickup-delay CLI option to reduce the number of TCP sessions that are synchronized by synchronizing TCP sessions only if they remain active for more than 30 seconds. Enabling this option could greatly reduce the number of sessions that are synchronized if a cluster typically processes very many short duration sessions, which is typical of most HTTP traffic for example.

Use the following command to enable a 30 second session pickup delay:

```
config system ha
  set session-pickup-delay enable
end
```

Enabling session pickup delay means that if a failover occurs more sessions may not be resumed after a failover. In most cases short duration sessions can be restarted with only a minor traffic interruption. However, if you notice too many sessions not resuming after a failover you might want to disable this setting.

The `session-pickup-delay` option applies to TCP sessions only and does not apply to connectionless and SCTP sessions.

FortiGate-6000 FGSP

FortiGate-6000 supports the FortiGate Session Life Support Protocol (FGSP) (also called standalone session sync) to synchronize sessions among up to four FortiGate-6000s. FortiGate-6000 also supports FGSP between FGCP clusters.

For details about FGSP, see: [FGSP](#).

You can have the following options for selecting interfaces to use for FGSP session synchronization:

- Up to eight physical data interfaces.
- One or more data interface LAGs.
- VLANs added to the data interfaces or data interface LAGs.
- The HA1 and HA2 interfaces.

FortiGate-6000 FGSP support has the following limitations:

- SLBC platforms do not support L2 session synchronization links. The `session-sync-dev` option is not supported.
- FortiGate-6000 FGSP doesn't support setting up IPv6 session filters using the `config session-sync-filter` option.
- Asymmetric IPv6 SCTP traffic sessions are not supported. These sessions are dropped.
- Inter-cluster session synchronization, or FGSP between FGCP clusters, is supported, see [Synchronizing sessions between FGCP clusters on page 82](#).
- FGSP IPsec tunnel synchronization is not supported.
- Fragmented packet synchronization is not supported.

You can use configuration synchronization to synchronize the configurations of the FortiGate-6000s in the FGSP deployment (see [Standalone configuration synchronization on page 87](#)). You can use the HA1 and HA2 interfaces for configuration synchronization. You can also configure the FortiGate-6000s separately or use FortiManager to keep key parts of the configuration, such as security policies, synchronized.

FGSP session synchronization options

FortiGate-6000 FGSP supports the following session synchronization options:

```
config system ha
  set session-pickup {disable | enable}
  set session-pickup-connectionless {disable | enable}
  set session-pickup-expectation {disable | enable}
  set session-pickup-nat {disable | enable}
  set session-pickup-delay {disable | enable}
end
```

Some notes:

- The `session-pickup-expectation` and `session-pickup-nat` options only apply to the FGSP. The FGCP synchronizes NAT sessions when you enable `session-pickup`.
- The `session-pickup-delay` option applies to TCP sessions only and does not apply to connectionless and SCTP sessions.
- The `session-pickup-delay` option should not be used in FGSP topologies where the traffic can take an asymmetric path (forward and reverse traffic going through different FortiGate-6000s).

Enabling session synchronization

Use the following command to synchronize TCP and SCTP sessions between FortiGate-6000s.

```
config system ha
  set session-pickup enable
end
```

Enabling `session-pickup` also enables session synchronization for connectionless protocol sessions, such as ICMP and UDP, by enabling `session-pickup-connectionless`. If you don't want to synchronize connectionless sessions, you can manually disable `session-pickup-connectionless`.

Synchronizing expectation sessions

Enable `session-pickup-expectation` to synchronize expectation sessions. FortiOS session helpers keep track of the communication of Layer-7 protocols such as FTP and SIP that have control sessions and expectation sessions. Usually the control sessions establish the link between server and client and negotiate the ports and protocols that will be used for data communications. The session helpers then create expectation sessions through the FortiGate for the ports and protocols negotiated by the control session.

The expectation sessions are usually the sessions that actually communicate data. For FTP, the expectation sessions transmit files being uploaded or downloaded. For SIP, the expectation sessions transmit voice and video data. Expectation sessions usually have a timeout value of 30 seconds. If the communication from the server is not initiated within 30 seconds the expectation session times out and traffic will be denied.

Synchronizing NAT sessions

Enable `session-pickup-nat` to synchronize NAT sessions.

Synchronizing sessions older than 30 seconds

Enable `session-pickup-delay` to synchronize TCP sessions only if they remain active for more than 30 seconds. This option improves performance when `session-pickup` is enabled by reducing the number of TCP sessions that are synchronized. This option does not affect SCTP or connectionless sessions.

Using data interfaces for FGSP session synchronization

FortiGate-6000 FGSP supports using up to eight physical data interfaces for FGSP session synchronization.

Use the following command to select up to eight physical data interfaces to use for FGSP session synchronization:

```
config system standalone-cluster
  set data-intf-session-sync-dev <interface-name> [<interface-name> ...]
end
```

You can use these individual interfaces or VLANs added to these interfaces for FGSP session synchronization. You can also create LAGs of two or more of these physical interfaces and use the LAGs for FGSP session synchronization. You can also add a VLAN to a LAG and use this VLAN for FGSP session synchronization.

Fortinet recommends:

- Use a data interface LAG for FGSP session synchronization. A LAG supports higher throughput than a single interface and also provides redundancy.
- Do not use FGSP session synchronization data interfaces for other traffic.
- Enable jumbo frames on the data interfaces, LAGs, and VLANs that you use for FGSP session synchronization.
- Keep the FGSP session synchronization data interfaces in a separate dedicated VDOM. Any VLANs you add to these interfaces or LAGs that you create for FGSP session synchronization should also be in the same dedicated VDOM. You must then specify this VDOM as the `peervd` in the `config system cluster-sync` configuration.

For example, you could create a VDOM called `fgsp-sync` and add the data interfaces, VLANs and LAGs that you are using for FGSP session synchronization to that VDOM. Then you can create the following `config cluster-peer` instance to synchronize sessions from the root VDOM:

```
config system standalone-cluster
  set data-intf-session-sync-dev <interface-name> [<interface-name> ...]
  config cluster-peer
    edit 1
      set peervd fgsp-sync
      set peerip <ip-address>
      set syncvd root
    end
  ...
end
```

Synchronizing sessions between FGCP clusters

FortiGate-6000 supports using FGSP to synchronize sessions among up to four FortiGate-6000 FGCP clusters. All of the FortiGate-6000s must be the same hardware model.

FGSP between FGCP clusters synchronizes sessions between the primary FortiGate-6000s in each cluster. FGCP HA then handles session synchronization between FortiGate-6000s in each FGCP cluster.

For details about FGSP between FGCP clusters, see: [Synchronizing sessions between FGCP clusters](#).

You can use data interfaces or data interface LAGs as FGSP session synchronization interfaces. The HA1 and HA2 interfaces are used for FGCP HA heartbeat between the FortiGate-6000s in each FGCP cluster.

FortiGate-6000 synchronizing sessions between FGCP clusters has the following limitations:

- The FGCP clusters cannot be configured for virtual clustering.
- NAT between the session synchronization interfaces is not supported.
- Standalone configuration synchronization between the FGCP clusters is not supported.
- Inter-cluster session synchronization doesn't support setting up IPv6 session filters using the `config session-sync-filter` option.
- When ICMP load balancing is set to `to-master`, ICMP packets are not installed on the DP processor. In an FGSP between FGCP session synchronization configuration with an asymmetry topology, synchronized ICMP packets will be dropped if the clusters have selected a different primary FPC. To avoid this possible traffic loss, set `dp-icmp-distribution-method` to `src-ip`, `dst-ip`, or `src-dst-ip`.
- Asymmetric IPv6 SCTP traffic sessions are not supported. These sessions are dropped.
- FGSP IPsec tunnel synchronization is not supported.
- Session synchronization packets cannot be fragmented. So the MTU for the session synchronization interface should be supported by the network.
- To reduce the number of failovers and the amount of session synchronization traffic, configuring HA override on the FGCP clusters is not recommended.

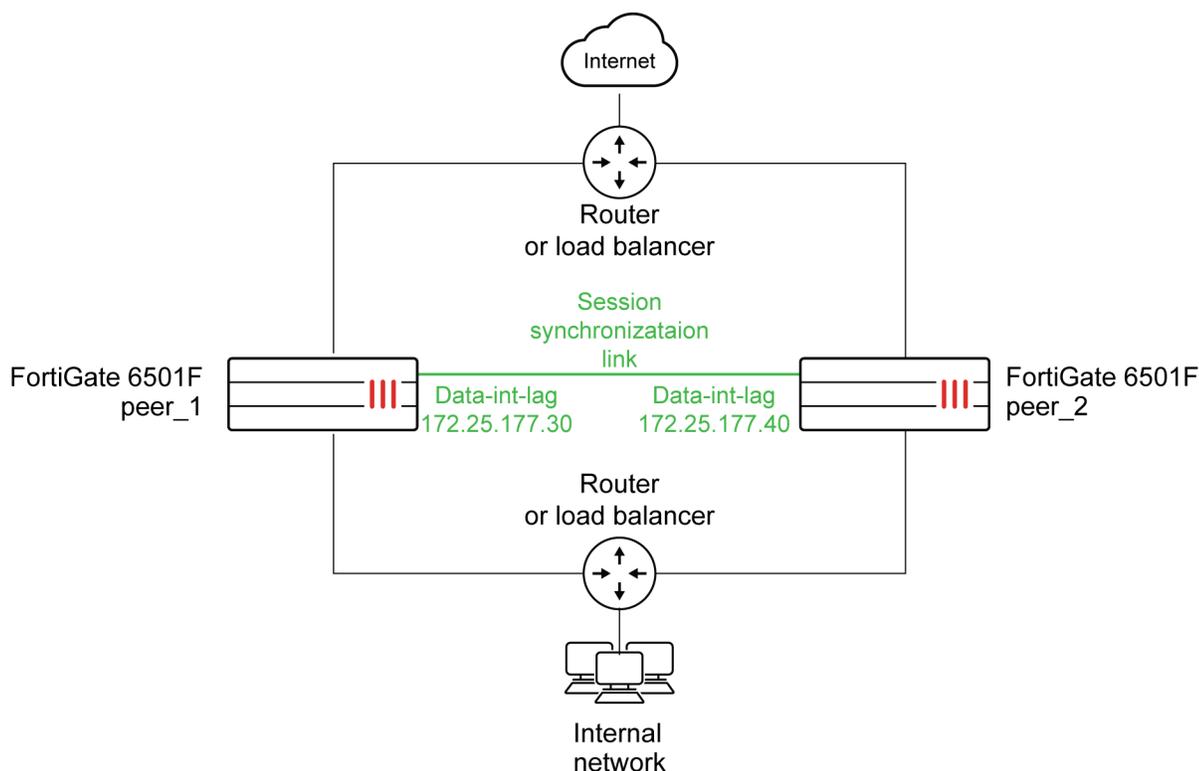
Example FortiGate-6000 FGSP session synchronization with a data interface LAG

This example shows how to configure FGSP to synchronize sessions between two FortiGate-6501s for the root VDOM and for a second VDOM, named `vdom-1`. For FGSP session synchronization, the example uses a data interface LAG that includes the port9 to port12 interfaces.

To set up the configuration, start by giving each FortiGate-6501F a different host name to make them easier to identify. This example uses `peer_1` and `peer_2`. On each FortiGate-6501F, create a VDOM named `fgsp-sync` and move the port9 to port12 interfaces to this VDOM. Then create a LAG named `Data-int-lag`, also in the `fgsp-sync` VDOM, that includes the port9 to port12 interfaces. The LAGs on both FortiGate-6501Fs are on the 172.25.177.0/24 network.

This example also adds standalone configuration synchronization using the HA1 and HA2 interfaces and sets the `peer_1` device priority higher so that it becomes the config sync primary. Once configuration synchronization is enabled, you can log into `peer_1` and add firewall policies and make other configuration changes and these configuration changes will be synchronized to `peer_2`. For information about configuration synchronization, including its limitations, see [Standalone configuration synchronization on page 87](#).

Example FortiGate-6000 FGSP configuration using data interface LAGs



1. Configure the routers or load balancers to distribute sessions to the two FortiGate-6501Fs.
2. Change the host names of the FortiGate-6501Fs to `peer_1` and `peer_2`.
3. Configure network settings for each FortiGate-6501F to allow them to connect to their networks and route traffic.
4. Add the `vdom-1` and `fgsp-sync` VDOMs to each FortiGate-6501F.
5. Also on each FortiGate-6501F, move the port9, port10, port11, and port12 data interfaces to the `fgsp-sync` VDOM.
6. On `peer_1`, configure the port9 to port12 interfaces to be FGSP session synchronization data interfaces.

```
config system standalone-cluster
  set standalone-group-id 6
  set group-member-id 1
  set data-intf-session-sync-dev port9 port10 port11 port12
end
```

7. On `peer_1`, add a data interface LAG to the `fgsp-sync` VDOM.

```
config system interface
  edit Data-int-lag
```

```

    set type aggregate
    set vdom fgsp-sync
    set member port9 port10 port11 port12
    set ip 172.25.177.30/24
    set mtu-override enable
    set mtu 9216
end

```

This configuration adds the data interface LAG to the fgsp-sync VDOM, includes the four data interfaces configured to be FGSP session synchronization interfaces, and configures the LAG to support jumbo frames.

8. On peer_1, configure session synchronization for the root and vdom-1 VDOMs.

```

config system standalone-cluster
  config cluster-peer
    edit 1
      set peervd fgsp-sync
      set peerip 172.25.177.40
      set syncvd root vdom-1
    end

```

peervd is fgsp-sync because the FGSP session synchronization data interfaces are in the fgsp-sync VDOM.

peerip is the IP address of the data interface LAG added to peer_2.

This configuration creates one cluster-sync instance that includes both VDOMs. You could have created a separate cluster-sync instance for each VDOM. If possible, however, avoid creating more than three cluster-sync instances. A fourth cluster-sync instance may experience reduced session synchronization performance.

9. On peer_1, enable configuration synchronization, enable session pickup, configure the heartbeat interfaces, and set a higher device priority. This makes peer_1 become the config sync primary.

```

config system ha
  set standalone-config-sync enable
  set session-pickup enable
  set session-pickup-connectionless enable
  set session-pickup-expectation enable
  set session-pickup-nat enable
  set priority 250
  set hbdev ha1 50 ha2 50
end

```

10. On peer_2, configure the port9 to port12 interfaces to be FGSP session synchronization data interfaces.

```

config system standalone-cluster
  set standalone-group-id 6
  set group-member-id 2
  set data-intf-session-sync-dev port9 port10 port11 port12
end

```

11. On peer_2, add a data interface LAG to the fgsp-sync VDOM:

```

config system interface
  edit Data-int-lag
    set type aggregate
    set vdom fgsp-sync
    set member port9 port10 port11 port12
    set ip 172.25.177.40/24
    set mtu-override enable
    set mtu 9216
  end

```

This configuration adds the data interface LAG to the fgsp-sync VDOM, includes the four data interfaces configured to be FGSP session synchronization interfaces, and configures the LAG to support jumbo frames.

12. On peer_2, configure session synchronization for the root and vdom-1 VDOMs.

```

config system standalone-cluster

```

```
config cluster-peer
  edit 1
    set peervd fgsp-sync
    set peerip 172.25.177.30
    set syncvd root vdom-1
  end
```

- 13.** On peer_2, enable configuration synchronization, enable session pickup, configure the heartbeat interfaces, and leave the device priority set to the default value.

```
config system ha
  set standalone-config-sync enable
  set session-pickup enable
  set session-pickup-connectionless enable
  set session-pickup-expectation enable
  set session-pickup-nat enable
  set hbdev ha1 50 ha2 50
end
```

As sessions are forwarded by the routers or load balancers to one of the FortiGate-6501Fs, the FGSP synchronizes the sessions to the other FortiGate-6501F. You can log into peer_1 and make configuration changes, which are synchronized to peer_2.

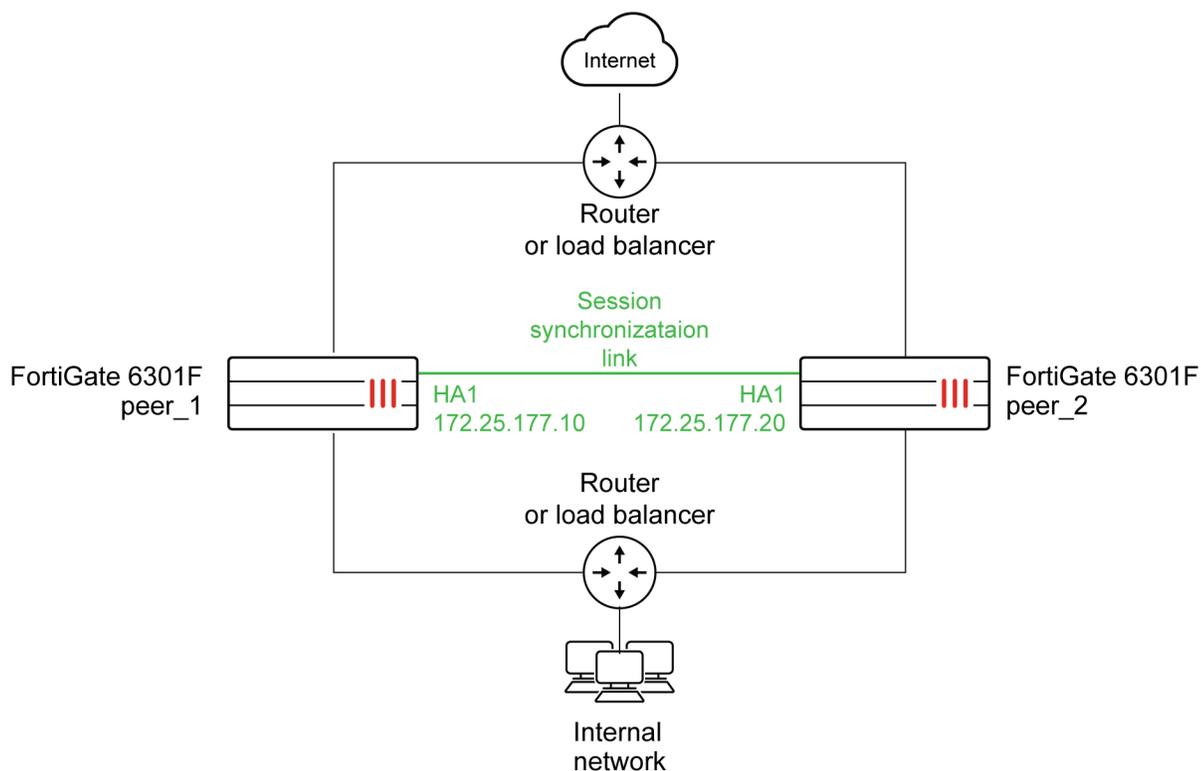
Example FortiGate-6000 FGSP session synchronization using HA interfaces

This example shows how to configure FGSP to synchronize sessions between two FortiGate-6301s for the root VDOM and for a second VDOM, named vdom-1. The example uses the HA1 interfaces of each FortiGate-6301F for session synchronization. The HA1 interfaces are connected to the 172.25.177.0/24 network. You could also connect and configure the HA2 interfaces and use them for session synchronization.

The interfaces of the two FortiGate-6301Fs must have their own IP addresses and their own network configuration. You can give the FortiGate-6301Fs different host names. This example uses peer_1 and peer_2, to make the FortiGate-6301Fs easier to identify.

This example also adds configuration synchronization and sets the peer_1 device priority higher so that it becomes the config sync primary. Once configuration synchronization is enabled, you can log into peer_1 and add firewall policies and make other configuration changes and these configuration changes will be synchronized to peer_2. For information about configuration synchronization, including its limitations, see [Standalone configuration synchronization on page 87](#).

Example FortiGate-6000 FGSP configuration



1. Configure the routers or load balancers to distribute sessions to the two FortiGate-6301Fs.
2. Change the host names of the FortiGate-6301Fs to peer_1 and peer_2.
3. Configure network settings for each FortiGate-6301F to allow them to connect to their networks and route traffic.
4. Add the vdom-1 VDOM to each FortiGate-6301F.
5. On peer_1, set up the standalone-cluster configuration.

```
config system standalone-cluster
  set standalone-group-id 3
  set group-member-id 1
end
```

6. On peer_1, configure the HA1 interface of with an IP address on the 172.25.177.0/24 network:

```
config system interface
  edit ha1
    set ip 172.25.177.10 255.255.255.0
  end
```

7. On peer_1, configure session synchronization for the root and vdom-1 VDOMs.

```
config system standalone-cluster
  config cluster-peer
    edit 1
      set peervd mgmt-vdom
      set peerip 172.25.177.20
      set syncvd root vdom-1
    end
```

Where, peervd will always be mgmt-vdom and peerip is the IP address of the HA1 interface of peer_2.

This configuration creates one cluster-sync instance that includes both VDOMs. You could have created a separate cluster-sync instance for each VDOM. If possible, however, avoid creating more than three cluster-

sync instances. A fourth `cluster-sync` instance may experience reduced session synchronization performance.

8. On `peer_1`, enable configuration synchronization, enable session pickup, configure the heartbeat interfaces, and set a higher device priority. This makes `peer_1` become the config sync primary.

```
config system ha
  set standalone-config-sync enable
  set session-pickup enable
  set session-pickup-connectionless enable
  set session-pickup-expectation enable
  set session-pickup-nat enable
  set priority 250
  set hbdev ha1 50 ha2 50
end
```

9. On `peer_2`, set up the standalone-cluster configuration.

```
config system standalone-cluster
  set standalone-group-id 3
  set group-member-id 2
end
```

10. On `peer_2`, configure the HA1 interface with an IP address on the 172.25.177.0/24 network:

```
config system interface
  edit ha1
    set ip 172.25.177.20 255.255.255.0
  end
```

11. On `peer_2`, configure session synchronization for the root and vdom-1 VDOMs.

```
config system standalone-cluster
  config cluster-peer
    edit 1
      set peervd mgmt-vdom
      set peerip 172.25.177.10
      set syncvd root vdom-1
    end
  end
```

12. On `peer_2`, enable configuration synchronization, enable session pickup, configure the heartbeat interfaces, and leave the device priority set to the default value.

```
config system ha
  set standalone-config-sync enable
  set session-pickup enable
  set session-pickup-connectionless enable
  set session-pickup-expectation enable
  set session-pickup-nat enable
  set hbdev ha1 50 ha2 50
end
```

As sessions are forwarded by the routers or load balancers to one of the FortiGate-6301Fs, the FGSP synchronizes the sessions to the other FortiGate-6301F. You can log into `peer_1` and make configuration changes, which are synchronized to `peer_2`.

Standalone configuration synchronization

FortiGate-6000 supports configuration synchronization (also called standalone configuration synchronization) for two FortiGate-6000s. Configuration synchronization means that most configuration changes made to one of the FortiGate-6000s are automatically synchronized to the other one.

For details about standalone configuration synchronization for FortiOS 6.0, see: [Standalone configuration sync](#).

Use the following command on both FortiGate-6000s to enable configuration synchronization:

```
config system ha
  set standalone-config-sync enable
end
```

In addition to enabling configuration synchronization, you must set up HA heartbeat connections between the FortiGate-6000s using the HA1 and HA2 interfaces. One HA heartbeat connection is required, two are recommended. Use the following command to enable heartbeat configuration for both the HA1 and HA2 interfaces. This command gives both heartbeat interfaces the same priority. You can choose to select different priorities for each heartbeat interface:

```
config system ha
  set hbdev ha1 50 ha2 50
end
```

When you enable configuration synchronization, configure and connect the heartbeat devices, FGCP primary unit selection criteria selects a config sync primary FortiGate-6000. Normally, the FortiGate-6000 with the highest serial number becomes the config sync primary and the other FortiGate-6000 becomes the config sync secondary.

All configuration changes that you make to the primary are synchronized to the secondary. To avoid synchronization problems, Fortinet recommends making all configuration changes to the primary.



See [Limitations on page 89](#) for a list of limitations of the configuration synchronization feature. Fortinet recommends disabling configuration synchronization once the configurations of the FortiGate-6000s have been synchronized.

Selecting the config sync primary FortiGate-6000

You can use device priority to select one of the FortiGate-6000s to become the config sync primary. For example, the following command enables configuration synchronization and sets a higher device priority than the default of 128 to make sure that this FortiGate-6000 becomes the primary.

```
config system ha
  set standalone-config-sync enable
  set priority 250
end
```

Settings that are not synchronized

Configuration synchronization does not synchronize settings that identify the FortiGate-6000 to the network. The following settings are not synchronized:

- Transparent mode management IPv4 and IPv6 IP addresses and default gateways.
- All `config system cluster-sync` settings.
- All `config system interface` settings except `vdom`, `vlanid`, `type` and `interface`.
- All `config firewall sniffer` settings.
- All router BFD and BFD6 settings.
- The following BGP settings: `as`, `router-id`, `aggregate-address`, `aggregate-address6`, `neighbor-group`, `neighbor`, `network`, and `network6`.
- The following OSPF settings: `router-id`, `area`, `ospf-interface`, `network`, `neighbor`, and `summary-address`.
- The following OSPF6 settings: `router-id`, `area`, and `ospf6-interface`.

- All RIP settings.
- All policy routing settings.
- All static routing settings.

Limitations

When configuration synchronization is enabled, there are some limitations, including but not limited to the following:

- Configuration synchronization does not support graceful HA firmware upgrades. If you upgrade the firmware of primary, the secondary also upgrades at the same time, disrupting network traffic. You can avoid traffic interruptions by disabling configuration synchronization and upgrading the firmware of each FortiGate-6000 separately.
- The configuration settings that are synchronized might not match your requirements. The current design and implementation of configuration synchronization is based on requirements from specific customers and might not work for your implementation.
- It can be difficult to control which FortiGate-6000 becomes the config sync primary and the config sync primary can dynamically change without notice. This could result in accidentally changing the configuration of the secondary or overwriting the configuration of the intended primary.

FortiGate-6000 VRRP HA

FortiGate-6000 supports the Virtual Router Redundancy Protocol (VRRP), allowing you to configure VRRP HA between FortiGate-6000 data interfaces. You can also add a FortiGate-6000 data interface to a VRRP domain with other VRRP routers.

To set up a FortiGate-6000 VRRP to provide HA for internet connectivity:

1. Add a virtual VRRP router to the internal interface to the FortiGate-6000(s) and routers to be in the VRRP domain.
2. Set the VRRP IP address of the domain to the internal network default gateway IP address.
3. Give one of the VRRP domain members the highest priority so it becomes the primary router and give the others lower priorities so they become secondary routers.

During normal operation, the primary VRRP router sends outgoing VRRP routing advertisements. Both the primary and backup VRRP routers listen for incoming VRRP advertisements from other routers in the VRRP domain. If the primary router fails, the new primary router takes over the role of the default gateway for the internal network and starts sending and receiving VRRP advertisements.

On the GUI you can go to **Network > Interfaces** and right click on the column header and add VRRP to the **Selected Columns** list to see the VRRP status of the data interfaces that are operating as VRRP routers.

For more information about FortiOS VRRP, see [FortiGate Handbook: VRRP](#).

Operating a FortiGate-6000

This chapter is a collection of information that you can use when operating your FortiGate-6000 system.

FortiLink support

FortiGate 6000F supports managing FortiSwitch devices over FortiLink. You can manage up to 300 FortiSwitch devices from one FortiGate 6000F.

Use the following command to enable Fortilink support on the GUI and in the CLI:

```
config system global
  set switch-controller enable
end
```

From the GUI you can go to **System > Feature Visibility** and enable **Switch Controller**.

Managed FortiSwitch GUI pages appear under the **WiFi & Switch Controller** GUI menu on all VDOMs except mgmt-vdom.

The default configuration includes a LAG named fortilink. To manage FortiSwitch devices you must add one or more interfaces to the fortilink LAG and connect managed FortiSwitch devices to these interfaces.

The FortiGate 6000F has the following FortiLink support limitation:

- FortiGate-6000 does not support upgrading managed FortiSwitch firmware from the **FortiOS Managed FortiSwitch GUI** page. Instead you must use the FortiGate-6000 CLI or log into the managed FortiSwitch to upgrade managed FortiSwitch firmware.

For more information about FortiLink support and managing FortiSwitches, see the [FortiSwitch Guide](#).

ECMP support

FortiGate-6000 supports most FortiOS IPv4 and IPv6 ECMP functionality. Before setting up an ECMP configuration you need to use the following command to configure the DP processor to operate with VDOM-based session tables:

```
config load-balance setting
  set dp-session-table-type vdom-based
end
```

Once you have enabled VDOM-based session tables, you can enable and configure ECMP as you would for any FortiGate.

VDOM-based session tables

In an ECMP configuration, because of load balancing, return traffic could enter through a different interface than the one it exited from. If this happens, the DP processor operating with default interface-based session tables may not be able to

send the return traffic to the FPC that processed the incoming session, causing the return traffic to be dropped. Operating with VDOM-based session tables solves this problem, allowing traffic received on a different interface to be properly identified and sent to the correct FPC.

Enabling VDOM session tables can reduce connections per second (CPS) performance so it should only be enabled if needed to support ECMP. This performance reduction can be more noticeable if the FortiGate-6000 is processing many firewall only sessions. If the FortiGate-6000 is performing content inspection where CPS performance is less important, the performance reduction resulting from enabling VDOM-based session tables may be less noticeable.

IPv4 and IPv6 ECMP load balancing

You can use the following command to configure the IPv4 ECMP load balancing method for a VDOM:

```
config system settings
  set v4-ecmp-mode {source-ip-based | weight-based | source-dest-ip-based}
end
```



With VDOM-based session tables enabled, the FortiGate 6000F supports all IPv4 ECMP load balancing methods supported by FortiOS except usage-based.

See this link for information about how to support IPv6 ECMP load balancing: [Technical Tip: ECMP – Load balancing algorithms for IPv4 and IPv6](#).

Enabling auxiliary session support

When ECMP is enabled, TCP traffic for the same session can exit and enter the FortiGate on different interfaces. To allow this traffic to pass through, FortiOS creates auxiliary sessions. Allowing the creation of auxiliary sessions is handed by the following command:

```
config system settings
  set auxiliary-sessions {disable | enable}
end
```

By default, the `auxiliary-session` option is disabled. This can block some TCP traffic when ECMP is enabled. If this occurs, enabling `auxiliary-session` may solve the problem. For more information, see [Technical Tip: Enabling auxiliary session with ECMP or SD-WAN](#).

ICAP support

You can configure your FortiGate-6000 to use Internet Content Adaptation Protocol (ICAP) to offload processing that would normally take place on the FortiGate-6000 to a separate server specifically set up for the required specialized processing.

ICAP servers are focused on a specific function, for example:

- Ad insertion
- Virus scanning
- Content translation

- HTTP header or URL manipulation
- Language translation
- Content filtering

FortiGate-6000 supports ICAP without any special configuration. This includes using ICAP to offload decrypted SSL traffic to an ICAP server. FortiOS decrypts the content stream before forwarding it to the ICAP server.

For more information about FortiOS support for ICAP, see [ICAP support](#).

Example ICAP configuration

ICAP is available for VDOMs operating in proxy mode. You can enable proxy mode from the **Global GUI** by going to **System > VDOM**, editing the VDOM for which to configure ICAP, and setting **Inspection Mode** to **Proxy**.

Then go to the VDOM, and go to **System > Feature Visibility** and enable **ICAP**.

From the CLI you can edit the VDOM, enable proxy inspection mode and enable ICAP. You can only enable ICAP from `config system settings` if proxy mode is already enabled.

```
config vdom
  edit VDOM-2
    config system settings
      set inspection-mode proxy
    end
  config system settings
    set gui-icap enable
  end
```

From the GUI you can add an ICAP profile by going to **Security Profiles > ICAP** and selecting **Create New** to create a new ICAP profile.

From the CLI you can use the following command to create an ICAP profile:

```
config icap profile
  edit "default"
  next
  edit "icap-test-profile"
    set request enable
    set response enable
    set request-server "icap-test"
    set response-server "icap-test"
    set request-failure bypass
    set response-failure bypass
    set request-path "echo"
    set response-path "echo"
  end
```

From the GUI you can add an ICAP server by going to **Security Profiles > ICAP Servers** and selecting **Create New** to create a new ICAP server.

From the CLI you can use the following command to create an ICAP server:

```
config icap server
  edit "icap-test"
    set ip-address 10.98.0.88
    set max-connections 1000
  end
```

Then create a firewall policy for the traffic to be sent to the ICAP server and include the ICAP profile.

```

config firewall policy
  edit 4
    set name "any-any"
    set uuid f4b612d0-2300-51e8-f15f-507d96056a96
    set srcintf "1-C1/5" "1-C1/6"
    set dstintf "1-C1/6" "1-C1/5"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ALL"
    set utm-status enable
    set logtraffic all
    set av-profile "default"
    set icap-profile "icap-test-profile"
    set profile-protocol-options "default"
    set ssl-ssh-profile "deep-inspection"
  end

```

SSL mirroring support

You can configure your FortiGate-6000 to "mirror" or send a copy of traffic decrypted by SSL inspection to one or more interfaces so that the traffic can be collected by a raw packet capture tool for archiving or analysis.



Decryption, storage, inspection, and use decrypted content is subject to local privacy rules. Use of these features could enable malicious users with administrative access to your FortiGate to harvest sensitive information submitted using an encrypted channel.

Use the information in [Mirroring SSL traffic in policies](#) to set up SSL mirroring for your FortiGate-6000.

You can use the following command from an FPC CLI to verify the mirrored traffic:

```

diagnose sniffer packet port20 'port 443' -c 50
interfaces=[port20]
filters=[port 443]
pcap_lookupnet: port20: no IPv4 address assigned
0.440714 8.1.1.69.18478 -> 9.2.1.130.443: syn 582300852
0.440729 9.2.1.130.443 -> 8.1.1.69.18478: syn 3198605956 ack 582300853
0.440733 8.1.1.69.18478 -> 9.2.1.130.443: ack 3198605957
0.440738 8.1.1.69.18478 -> 9.2.1.130.443: psh 582300853 ack 3198605957
0.441450 9.2.1.130.443 -> 8.1.1.69.18478: psh 3198605957 ack 582301211
0.441535 9.2.1.130.443 -> 8.1.1.69.18478: psh 3198607351 ack 582301211
0.441597 9.2.1.130.443 -> 8.1.1.69.18478: psh 3198608747 ack 582301211
0.441636 9.2.1.130.443 -> 8.1.1.69.18478: psh 3198610143 ack 582301211
0.441664 9.2.1.130.443 -> 8.1.1.69.18478: psh 3198611539 ack 582301211
0.441689 9.2.1.130.443 -> 8.1.1.69.18478: psh 3198612935 ack 582301211
0.441715 9.2.1.130.443 -> 8.1.1.69.18478: psh 3198614331 ack 582301211
0.441739 9.2.1.130.443 -> 8.1.1.69.18478: psh 3198615727 ack 582301211
0.441764 9.2.1.130.443 -> 8.1.1.69.18478: psh 3198617123 ack 582301211

```

VXLAN support

FortiGate 6000F supports terminating VXLAN traffic using VXLAN interfaces. VXLAN traffic cannot be load balanced, so you should use a flow rule similar to the following to send all VXLAN traffic terminated by the FortiGate 6000F to the primary FPC:

```
config load-balance flow-rule
  edit 0
    set status enable
    set ether-type ip
    set protocol 17
    set forward-slot master
    set src-interface <local LAN>
    set dst-l4port 4789-4789
    set comment "vxlan"
  end
```

`dst-l4port` must be set to the VXLAN destination port. The default VXLAN destination port is 4789. You should change the port number range in the flow rule if you change the VXLAN port number.

FortiGate 6000F IPsec load balancing EMAC VLAN interface limitation

On a FortiGate 6000F, because of a DP processor limitation, IPsec VPN load balancing is not supported for sessions received by an EMAC VLAN interface that is not in the same VDOM as the interface that the EMAC VLAN interface has been added to.

The following workarounds are available:

- Change the FortiGate 6000F configuration so that the EMAC VLAN interface is in the same VDOM as the interface that the EMAC VLAN interface is added to (the EMAC VLAN interface is in the same VDOM as its parent interface).
- Disable IPsec VPN load balancing and configure the IPsec phase 1 to send packets to the primary FPC or to a specific FPC. If you have multiple IPsec VPNs, you can achieve some load balancing by configuring different IPsec phase 1 configurations to send packets to different FPCs.

In addition, for each IPsec phase 1, create a flow rule to forward clear-text traffic from the EMAC VLAN interface to the primary FPC or to a specific FPC. The FPC in the flow rule must match the FPC in the IPsec phase 1 configuration.

- Do not use EMAC VLAN interfaces. For example, you could use standard VLAN interfaces. This may require using an external switch to handle VLAN tagging.

Configuring a FortiGate-6000 to operate in FIPS-CC mode

If the version of FortiOS running on your FortiGate 6000F supports FIPS-CC mode, you can use the `config system fips-cc` command to switch your FortiGate 6000F to operate in FIPS-CC mode.

When you enter this command on most FortiGate models, the FortiGate restarts, generates a new set of encryption keys, runs a set of startup and conditional self-tests, and then starts up operating in FIPS-CC mode.

The FortiGate 6000F follows the same process except that first the management board and then all of the FPCs each generate their own sets of keys and then run their own set of startup and conditional self tests.

To make sure the conversion goes smoothly, you should make sure all of the FPCs are synchronized with the management board before switching to FIPS-CC mode. From the management board CLI, you can run the `diagnose load-balance status` command to confirm that the `Status Message` for all FPCs is `Running`.

```
diagnose load-balance status
=====
MBD SN: F6KF31T018900143
  Primary FPC Blade: slot-1

Slot 1: FPC6KFT018901327
  Status:Working  Function:Active
  Link:      Base: Up      Fabric: Up
  Heartbeat: Management: Good  Data: Good
  Status Message:"Running"
Slot 2: FPC6KFT018901372
  Status:Working  Function:Active
  Link:      Base: Up      Fabric: Up
  Heartbeat: Management: Good  Data: Good
  Status Message:"Running"
Slot 3: FPC6KFT018901346
  Status:Working  Function:Active
  Link:      Base: Up      Fabric: Up
  Heartbeat: Management: Good  Data: Good
  Status Message:"Running"
Slot 4: FPC6KFT018901574
  Status:Working  Function:Active
  Link:      Base: Up      Fabric: Up
  Heartbeat: Management: Good  Data: Good
  Status Message:"Running"
Slot 5: FPC6KFT018901345
  Status:Working  Function:Active
  Link:      Base: Up      Fabric: Up
  Heartbeat: Management: Good  Data: Good
  Status Message:"Running"
Slot 6: FPC6KFT018901556
  Status:Working  Function:Active
  Link:      Base: Up      Fabric: Up
  Heartbeat: Management: Good  Data: Good
  Status Message:"Running"
```



If one or more FPCs are not running or are not synchronized you may need to wait a bit longer for the FPC to start up and become synchronized. You could also try manually restarting the FPC or check [Troubleshooting an FPC failure on page 108](#) for other suggestions.

If all of the FPCs are running and synchronized, you can enter the `config system fips-cc` command from the management board CLI. If you are logged into the management board CLI using a console connection, messages similar to the following appear as the management board completes its self tests and then waits for the FPCs to complete their self-tests:

```
FIPS-CC mode: Starting self-tests.  
Running Configuration/VPN Bypass test... passed  
Running AES test... passed  
Running SHA1 HMAC test... passed  
Running SHA256 HMAC test... passed  
Running SHA384/512 HMAC test... passed  
Running RSA test... passed  
Running ECDSA test... passed  
Running TLS1.1-KDF test... passed  
Running TLS1.2-KDF test... passed  
Running SSH-KDF test... passed  
Running IKEv1-KDF test... passed  
Running IKEv2-KDF test... passed  
Running Primitive-Z test... passed  
Running Firmware integrity test... passed  
Running RBG-instantiate test... passed  
Running RBG-reseed test... passed  
Running RBG-generate test... passed  
Motherboard Self-tests passed  
Please wait for FPC self-tests to complete
```

As each FPC completes its self-tests, the FPC sends the results (pass or fail) to the management board. Each FPC also records log messages with the self-test results. Until all of the FPCs have reported successfully passing their self-tests, the front panel interfaces remain down.

This may take a few minutes. When all of the FPCs pass their self-tests, the following message appears on the management board console connection:

```
FPC self-tests have completed
```

The login prompt appears and you can log into the management board CLI. The front panel interfaces come up as well.

You can use the `get system status` command to verify that the FortiGate 6000F is operating in FIPS-CC mode.

```
get system status  
. . .  
FIPS-CC mode: enable  
. . .
```

Troubleshooting FortiGate 6000F self tests

Since the management board and all of the FPCs have to pass their self-tests, converting a FortiGate 6000F to FIPS-CC mode may take longer than expected and may be more prone to temporary failure than expected.

All FPCs must pass their self-tests before the self-test timer expires. So if the timer is set to 1440 seconds, the management board will wait up to 1440 seconds to receive self-test pass messages from all of the FPCs. If the self-test timer expires before all of the FPCs pass their self-tests, the FortiGate 6000F keeps running but all interfaces remain shut down.

The self-test timer gives you time to check the status of the FPCs and troubleshoot and resolve any problems that may prevent them from starting up or passing their self-tests.

Fortinet recommends that once the management board passes its self tests, run the `diagnose load-balance status` command to confirm that the Status Message for all FPCs is Running.

```
diagnose load-balance status
=====
MBD SN: F6KF31T018900143
  Primary FPC Blade: slot-1

  Slot 1: FPC6KFT018901327
    Status:Working  Function:Active
    Link:           Base: Up           Fabric: Up
    Heartbeat: Management: Good  Data: Good
    Status Message:"Running"
  Slot 2: FPC6KFT018901372
    Status:Working  Function:Active
    Link:           Base: Up           Fabric: Up
    Heartbeat: Management: Good  Data: Good
    Status Message:"Running"
  Slot 3: FPC6KFT018901346
    Status:Working  Function:Active
    Link:           Base: Up           Fabric: Up
    Heartbeat: Management: Good  Data: Good
    Status Message:"Running"
  Slot 4: FPC6KFT018901574
    Status:Working  Function:Active
    Link:           Base: Up           Fabric: Up
    Heartbeat: Management: Good  Data: Good
    Status Message:"Running"
  Slot 5: FPC6KFT018901345
    Status:Working  Function:Active
    Link:           Base: Up           Fabric: Up
    Heartbeat: Management: Good  Data: Good
    Status Message:"Running"
  Slot 6: FPC6KFT018901556
    Status:Working  Function:Active
    Link:           Base: Up           Fabric: Up
    Heartbeat: Management: Good  Data: Good
    Status Message:"Running"
```

If the FortiGate 6000F has just started up, some of the FPCs may not be in the running state because they are still starting up. Try running the `diagnose load-balance status` command a few more times to see if all of the FPCs transition to the running state.

If an FPC continues to not be in the running state, you can try manually restarting it. You can also use the information in [Troubleshooting an FPC failure on page 108](#) to do further investigation.

If an FPC fails its self-test, the management board console may display a message similar to the following (which indicates that the FPC in slot 2 experienced a self-test failure):

```
Self-test failure: FPC 0000002
```

The FPC may self-correct and re-try and pass the self-test without any intervention. You could also try manually restarting the FPC or check [Troubleshooting an FPC failure on page 108](#) for other suggestions.

Global option for proxy-based certificate queries

In some cases you may want to be able to send certificate queries using a FortiGate 6000F management interface instead of a data interface. FortiGate 6000F includes the following global command that you can use to enable or disable using a data interface or a system management interface for certificate queries for proxy-based firewall policies.

```
config global
  config system global
    set proxy-cert-use-mgmt-vdom {disable | enable}
  end
```

This option is disabled by default and by default data interfaces are used to send certificate queries for proxy-based firewall policies. Enable this option to send certificate queries for proxy-based firewall policies through the mgmt-vdom VDOM using FortiGate-6000 management interfaces.

Using data interfaces for management traffic

You can set up IPv4 and IPv6 in-band management connections to all FortiGate 6000F data interfaces by setting up administrative access for the data interface that you want to use to manage the FortiGate 6000F. For in-band management of a transparent mode VDOM, you must also set up the transparent mode management IP address.

Connecting to a data interface for management is the same as connecting to one of the management interfaces. For example, you can log in to the GUI or CLI of the FortiGate 6000F management board.

Administrators with VDOM-level access can log into to their VDOM if they connect to a data interface that is in their VDOM.

In-band management limitations

In-band management has the following limitations:

- In-band management does not support using special port numbers to connect to individual FPCs or the management board. If you have logged in using an in-band management connection, the special management HTTPS port numbers appear on the Security Fabric dashboard widget when you hover over individual FPCs. You can click on an FPC in the Security Fabric dashboard widget and select **Login to...** to log into the GUI of that FPC. This action creates an out-of-band management connection by crafting a URL that includes the IP address of the FortiGate 6000F mgmt1 plus the special HTTPS port number required to connect to that FPC.
- SNMP in-band management is not supported.
- VRF routes are not applied to outgoing in-band management traffic.
- Changes made on the fly to administrative access settings are not enforced for in-progress in-band management sessions. The changes apply to new in-band sessions only. For example, if an administrator is using SSH for an in-band management connection and you change the SSH administrative port, that in-band management session can continue. Any out-of-band management sessions would need to be restarted with the new port number. New in-band SSH management sessions need to use the new port number. HTTPS access works the same way, however, HTTPS starts new sessions every time you navigate to a new GUI page. So an on the fly change would affect an HTTPS in-band management session whenever the administrator navigates to a new GUI page.

FortiGate 6000F management interface LAG and VLAN support

FortiGate 6000F supports adding the mgmt1 and mgmt2 interfaces to an LACP link aggregation group (LAG). You can also add VLAN interfaces to the mgmt1, mgmt2, and mgmt3 interfaces or to a LAG that includes mgmt1 and mgmt2.

You can use the following configuration to create a management interface LAG that includes the mgmt1 and mgmt2 interfaces.

```
config system interface
  edit "lacp_mgmt"
    set vdom mgmt-vdom
    set type aggregate
    set member mgmt1 mgmt2
  end
```



To be able to add an interface to a LAG you must remove all references to that interface (including static routes) and unset the IP address of the interface.

The management interface LAG fully supports LACP and supports other standard interface features. The management interface LAG as well as any VLAN interfaces added to the mgmt1, mgmt2, or mgmt3 interfaces or to the management interface LAG must remain in the mgmt-vdom VDOM.

Management interface LAG limitations

Management interface LAG support has the following limitations:

- You cannot set a management interface LAG to be the SLBC management interface by adding it to the `config load-balance setting slbc-mgmt-intf` option. This means that you cannot use the management interface LAG IP address with special port numbers to access the management board or individual FPCs as described in [Special management port numbers on page 27](#).
After creating a management interface LAG, if you still want to be able to use special port numbers to log into the management board or individual FPCs, you can use the mgmt3 interface for this access by setting `slbc-mgmt-intf` to `mgmt3` and connecting MGMT3 to the management network.
- FPCs and the management board assign different MAC addresses to the management interface LAG. The management board uses the MAC address of the second interface in the member list while the FPCs use the MAC address of the first interface in the member list.
- You can add the mgmt3 interface to the same LAG as mgmt1 and mgmt2. This configuration is not recommended, since LACP may not work as expected if the LACP group contains interfaces with different speeds. Adding mgmt3 might work in some configurations.
- You can add mgmt1, mgmt2, or mgmt3 to a LAG even if the management interface is configured as the SLBC management interface.
- If mgmt1, mgmt2, or mgmt3 are HA monitored interfaces they cannot be added to a management interface LAG.

Setting the MTU for a data interface

You can use the following command to change the MTU for a FortiGate-6000 data interface:

```
config system interface
edit port10
set mtu-override enable
set mtu <value>
end
```

For the FortiGate 6000F the default <value> is 1500 and the range is 256 to 9216.

More management connections than expected for one device

The FortiGate-6000 may show more management-related network activity than most FortiGate devices. This occurs because many management functions are handled independently by the management board and each FPC.

For example, when a FortiGate-6000 first starts up, the management board and all of the FPCs perform their DNS lookups. Resulting in more DNS-related traffic during startup than expected for a single device. Once the system is processing data traffic, the amount of management traffic would be proportional to the amount of traffic the system is processing.

More ARP queries than expected for one device - potential issue on large WiFi networks

The FortiGate-6000 sends more ARP queries than expected because each FPC builds its own ARP table to be able to communicate with devices in the same broadcast domain or layer 2 network. (This also may result in FPCs and the management board all having different ARP tables.) This behavior does not cause a problem with most layer 2 networks. However, because the ARP traffic for all of the FPCs comes from the same mac and IP address, on networks with broadcast filtering or ARP suppression, some of the FortiGate-6000 ARP queries and replies may be suppressed. If this happens, FPCs may not be able to build complete ARP tables. An FPC with an incomplete ARP table will not be able to forward sessions to some destinations that it should be able to reach, resulting in dropped sessions.

Broadcast filtering or ARP suppression is commonly used on large WiFi networks to control the amount of ARP traffic on the WiFi network. Dropped FortiGate-6000 sessions have been seen when a FortiGate-6000 is connected to the same broadcast domain as a large WiFi network with ARP suppression.

To resolve this dropped session issue, you can remove broadcast filtering or ARP suppression from the network. If this is not an option, Fortinet recommends that you install a layer 3 device to separate the FortiGate-6000 from the WiFi network broadcast domain. ARP traffic is reduced because the FPCs no longer need to add the addresses of all of the WiFi devices to their ARP tables since they are on a different broadcast domain. The FPCs just need to add the address of the layer 3 device.

VLAN ID 1 is reserved

When setting up VLANs, do not set the VLAN ID to 1. This VLAN ID is reserved by FortiOS. Any configurations that use a VLAN with VLAN ID = 1 will not work as expected.

Connecting to FPC CLIs using the console port

If you connect a PC to the FortiGate-6000 console port with a serial cable and open a terminal session, you are connected to the management board CLI. You can press Ctrl-T to enable console switching mode. Pressing Ctrl-T multiple times cycles through the management board (MBD) CLI and FPC CLIs. Once you have connected to the CLI that you want to use, press Enter to enable the CLI and log in.

The default settings for connecting to the console port are:

Baud Rate (bps) 9600, Data bits 8, Parity None, Stop bits 1, and Flow Control None.

Firmware upgrade basics

The management board and the FPCs in your FortiGate 6000F system run the same firmware image. You upgrade the firmware from the management board GUI or CLI just as you would any FortiGate product.

You can perform a graceful firmware upgrade of an FGCP cluster by setting `upgrade-mode` to `uninterruptible` and enabling `session-pickup`. A graceful firmware upgrade only causes minimal traffic interruption.

Upgrading the firmware of a standalone FortiGate 6000F, or FortiGate 6000F HA cluster with `upgrade-mode` set to `simultaneous` interrupts traffic because the firmware running on the management board and all of the FPCs upgrades in one step. These firmware upgrades should be done during a quiet time because traffic will be interrupted during the upgrade process.

A firmware upgrade takes a few minutes, depending on the number of FPCs in your FortiGate 6000F system. Some firmware upgrades may take longer depending on factors such as the size of the configuration and whether an upgrade of the DP3 processor is included.

Before beginning a firmware upgrade, Fortinet recommends that you perform the following tasks:

- Review the latest release notes for the firmware version that you are upgrading to.
- Verify the recommended upgrade path, as documented in the release notes.
- Back up your FortiGate 6000F configuration.



To make sure a FortiGate 6000F firmware upgrade is successful, before starting the upgrade Fortinet recommends you use health checking to make sure the management board and the FPCs are all synchronized and operating as expected.

If you are following a multi-step upgrade path, you should re-do health checking after each upgrade step to make sure all components are synchronized before the next step.

You should also perform a final round of health checking after the firmware upgrade process is complete.

For recommended health checking commands, see the following Fortinet community article:

[Technical Tip: FortiGate-6000/7000 Chassis health check commands.](#)



Fortinet recommends that you review the services provided by your FortiGate 6000F before a firmware upgrade and then again after the upgrade to make sure that these services continue to operate normally. For example, you might want to verify that you can successfully access an important server used by your organization before the upgrade and make sure that you can still reach the server after the upgrade and performance is comparable. You can also take a snapshot of key performance indicators (for example, number of sessions, CPU usage, and memory usage) before the upgrade and verify that you see comparable performance after the upgrade.

Installing firmware on an individual FPC

You may want to install firmware on an individual FPC to resolve a software-related problem with the FPC or if the FPC is not running the same firmware version as the management board. The following procedure describes how to transfer a new firmware image file to the FortiGate 6000F internal TFTP server and then install the firmware on an FPC.

1. Copy the firmware image file to a TFTP server, FTP server, or USB key.
2. To upload the firmware image file onto the FortiGate 6000F internal TFTP server, from the management board CLI, enter one of the following commands.

- To upload the firmware image file from an FTP server:

```
execute upload image ftp <image-file-and-path> <comment> <ftp-server-address>
<username> <password>
```

- To upload the firmware image file from a TFTP server:

```
execute upload image tftp <image-file> <comment> <tftp-server-address>
```

- To upload the firmware image file from a USB key plugged into the FortiGate 6000F USB port:

```
execute upload image usb <image-file-and-path> <comment>
```

3. Enter the following command to confirm that the firmware image is available on the internal tftp server.

```
fnsysctl ls data2/tftproot
```

This command lists the files and folders available in the `tftproot` directory, one of them should be the image file that you uploaded to the TFTP server in the previous step.

4. Enter the following command to check the IP address of the internal TFTP server:

```
fnsysctl ifconfig base-tftp
base-tftp Link encap:Ethernet HWaddr 02:1C:BA:54:92:9B
inet addr:169.254.255.1 Bcast:169.254.255.255 Mask:255.255.255.0
UP BROADCAST RUNNING MULTICAST MTU:1492 Metric:1
RX packets:346263 errors:0 dropped:0 overruns:0 frame:0
TX packets:346193 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:17315539 (16.5 MB) TX bytes:270658277 (258.1 MB)
```

The IP address in the `inet addr` field is the IP address of the internal TFTP server and should be 169.254.255.1.

5. Use the `execute system console-server` command to log into to the console of the FPC that you want to install the new firmware on.

For example, to log into the console of the FPC in slot 3 enter

```
execute system console-server connect 3
```

Log into the console using your FortiGate 6000F administrator name and password.

- Restart the FPC. You can do this from the FPC console by accessing the global configuration and entering the `execute reboot` command:

```
config system global
  execute reboot
```

- When the FPC starts up, follow the boot process in the console session, and press any key when prompted to interrupt the boot process.
- To set up the TFTP configuration, press C.
- Use the BIOS menu to set the following. Change settings only if required.

```
[P]: Set image download port: MGMT1
```

```
[D]: Set DHCP mode: Disabled
```

```
[I]: Set local IP address: Enter an IP address on the same network as the TFTP sever IP address. For example, if the TFTP server IP address is 169.254.255.1, set the local IP address to 169.254.255.2.
```

```
[S]: Set local Subnet Mask: 255.255.255.0
```

```
[G]: Set local gateway: 169.254.255.1 (same as the TFTP server IP address).
```

```
[V]: Local VLAN ID: Should be set to <none>. (use -1 to set the Local VLAN ID to <none>.)
```

```
[T]: Set remote TFTP server IP address: 169.254.255.1 (The IP address of the internal TFTP server.)
```

```
[F]: Set firmware image file name: The name of the firmware image file uploaded to the internal TFTP server.
```

- To quit this menu, press Q.
- To review the configuration, press R.
To make corrections, press C and make the changes as required.
When the configuration is correct, proceed to the next step.
- To start the TFTP transfer, press T.
The FPC downloads the firmware image from the internal TFTP server and installs it. The FPC then restarts with its configuration reset to factory defaults. The configuration is then synchronized from the management board to the FPC. The FPC restarts again, rejoins the cluster, and is ready to process traffic.
- To verify that the configuration of the FPC has been synchronized, enter the `diagnose sys confsync status | grep in_sy` command. The command output below shows an example of the synchronization status of some of the FPCs in an HA cluster of two FortiGate 6301F devices. The field `in_sync=1` indicates that the configuration of the FPC is synchronized.

```
FPC6KFT018901327, Secondary, uptime=615368.33, priority=19, slot_id=1:1, idx=1, flag=0x4, in_sync=1
F6KF31T018900143, Primary, uptime=615425.84, priority=1, slot_id=1:0, idx=0, flag=0x10, in_sync=1
FPC6KFT018901372, Secondary, uptime=615319.63, priority=20, slot_id=1:2, idx=1, flag=0x4, in_sync=1
F6KF31T018900143, Primary, uptime=615425.84, priority=1, slot_id=1:0, idx=0, flag=0x10, in_sync=1
FPC6KFT018901346, Secondary, uptime=423.91, priority=21, slot_id=1:3, idx=1, flag=0x4, in_sync=1
```

FPCs that are missing or that show `in_sync=0` are not synchronized. To synchronize an FPC that is not synchronized, log into the CLI of the FPC and restart it using the `execute reboot` command. If this does not solve the problem, contact Fortinet Support at <https://support.fortinet.com>.

The example output also shows that the uptime of the FPC in slot 3 is lower than the uptime of the other FPCs, indicating that the FPC in slot 3 has recently restarted.

If you enter the `diagnose sys confsync status | grep in_sy` command before an FPC has completely restarted, it will not appear in the output. Also, the Cluster Status dashboard widget will temporarily show that it is not synchronized.

Installing firmware from the BIOS after a reboot

A common method for resetting the configuration of a FortiGate involves installing firmware by restarting the FortiGate, interrupting the boot process, and using BIOS prompts to download a firmware image from a TFTP server. This process is also considered the best way to reset the configuration of your FortiGate.



Installing or upgrading FortiGate 6000F firmware from the BIOS after a reboot installs firmware on and resets the configuration of the management board only. FPCs will continue to operate with their current configuration and firmware build. The FortiGate-6000 system does not synchronize firmware upgrades that are performed from the BIOS. After you install firmware on the management board from the BIOS after a reboot, you must synchronize the new firmware build and configuration to the FPCs.

Installing or upgrading FortiGate 6001F, FortiGate 6501F, or FortiGate 6301F firmware from the BIOS after a reboot disables the log disk RAID configuration. You must rebuild the RAID configuration for normal log disk operation. If the FortiGate 6001F, FortiGate 6501F, or FortiGate 6301F is part of an FGCP HA cluster, both FortiGates in the cluster must have the same log disk RAID configuration. See [Changing the FortiGate-6301F and 6501F log disk and RAID configuration on page 111](#).

Use the following steps to upload firmware from a TFTP server to the management board. This procedure involves creating a connection between the TFTP server and one of the MGMT interfaces.

This procedure also involves connecting to the management board CLI using the FortiGate 6000F console port, rebooting the management board, interrupting the boot from the console session, and following BIOS prompts to install the firmware. During this procedure, the FortiGate 6000F will not be able to process traffic.

1. Set up a TFTP server and copy the firmware file to the TFTP server default folder.
2. Set up your network to allow traffic between the TFTP server and one of the management interfaces, (for example, MGMT1).
3. Using the console cable supplied with your FortiGate 6000F, connect the console port on the FortiGate to a USB port on your management computer.
4. Start a terminal emulation program on the management computer. Use these settings: Baud Rate (bps) 9600, Data bits 8, Parity None, Stop bits 1, and Flow Control None.
5. Log in to the management board CLI.
6. To restart the management board, enter the `execute reboot` command.
7. When the management board starts up, follow the boot process in the terminal session, and press any key when prompted to interrupt the boot process.
8. To set up the TFTP configuration, press C.
9. Use the BIOS menu to set the following. Change settings only if required.

[P]: Set image download port: MGMT1 (the connected MGMT interface)

[D]: Set DHCP mode: Disabled

[I]: Set local IP address: The IP address of the MGMT interface that you want to use to connect to the TFTP server. This address can be the same as the FortiGate 6000F management IP address and cannot conflict with other addresses on your network.

[S]: Set local Subnet Mask: Set as required for your network.

[G]: Set local gateway: Set as required for your network.

[V]: Local VLAN ID: Should be set to <none>. (use -1 to set the Local VLAN ID to <none>.)

[T]: Set remote TFTP server IP address: The IP address of the TFTP server.

[F]: Set firmware image file name: The name of the firmware image file that you want to install.

10. To quit this menu, press Q.
11. To review the configuration, press R.
To make corrections, press C and make the changes as required. When the configuration is correct, proceed to the next step.
12. To start the TFTP transfer, press T.
The management board downloads the firmware image from the TFTP server and installs it on the management board. The management board then restarts with its configuration reset to factory defaults.
13. Once the management board restarts, verify that the correct firmware is installed.
You can do this from the management board GUI dashboard or from the CLI using the `get system status` command.



If you are installing firmware on a FortiGate 6001F, FortiGate 6501F, or FortiGate 6301F, the log disk RAID configuration will be disabled once the management board restarts. You must rebuild the RAID configuration for normal log disk operation. If the FortiGate 6001F, FortiGate 6501F, or FortiGate 6301F is part of an FGCP HA cluster, both FortiGates in the cluster must have the same log disk RAID configuration. See [Changing the FortiGate-6301F and 6501F log disk and RAID configuration on page 111](#).

14. Continue by [Synchronizing the FPCs with the management board on page 105](#).

Synchronizing the FPCs with the management board

After you install firmware on the management board from the BIOS after a reboot, the firmware version and configuration of the management board will most likely not be synchronized with the FPCs. You can verify this from the management board CLI using the `diagnose sys confsync status | grep in_sy` command. The `in_sync=0` entries in the following example output for a FortiGate 6301F show that the management board (serial number ending in 143) is not synchronized with the FPCs.

```
diagnose sys confsync status | grep in_sy
FPC6KFT018901327, Secondary, uptime=59.44, priority=19, slot_id=1:1, idx=1, flag=0x4, in_sync=0
F6KF31T018900143, Primary, uptime=119.72, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
FPC6KFT018901372, Secondary, uptime=58.48, priority=20, slot_id=1:2, idx=1, flag=0x4, in_sync=0
F6KF31T018900143, Primary, uptime=119.72, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
FPC6KFT018901346, Secondary, uptime=58.44, priority=21, slot_id=1:3, idx=1, flag=0x4, in_sync=0
F6KF31T018900143, Primary, uptime=119.72, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
FPC6KFT018901574, Secondary, uptime=58.43, priority=22, slot_id=1:4, idx=1, flag=0x4, in_sync=0
F6KF31T018900143, Primary, uptime=119.72, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
FPC6KFT018901345, Secondary, uptime=57.40, priority=23, slot_id=1:5, idx=1, flag=0x4, in_sync=0
F6KF31T018900143, Primary, uptime=119.72, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
FPC6KFT018901556, Secondary, uptime=58.43, priority=24, slot_id=1:6, idx=1, flag=0x4, in_sync=0
F6KF31T018900143, Primary, uptime=119.72, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
F6KF31T018900143, Primary, uptime=119.72, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
FPC6KFT018901327, Secondary, uptime=59.44, priority=19, slot_id=1:1, idx=1, flag=0x4, in_sync=0
FPC6KFT018901345, Secondary, uptime=57.40, priority=23, slot_id=1:5, idx=2, flag=0x4, in_sync=0
FPC6KFT018901346, Secondary, uptime=58.44, priority=21, slot_id=1:3, idx=3, flag=0x4, in_sync=0
FPC6KFT018901372, Secondary, uptime=58.48, priority=20, slot_id=1:2, idx=4, flag=0x4, in_sync=0
FPC6KFT018901556, Secondary, uptime=58.43, priority=24, slot_id=1:6, idx=5, flag=0x4, in_sync=0
FPC6KFT018901574, Secondary, uptime=58.43, priority=22, slot_id=1:4, idx=6, flag=0x4, in_sync=0
```

You can also verify the synchronization status from the management board Configuration Sync Monitor.

To re-synchronize the FortiGate 6000F, which has the effect of resetting all of the FPCs, re-install firmware on the management board.



You can also manually install firmware on each FPC from the BIOS after a reboot. This multi-step manual process is just as effective as installing the firmware for a second time on the management board to trigger synchronization to the FPCs, but takes much longer.

1. Log in to the management board GUI.
2. Install a firmware build on the management board from the GUI or CLI. The firmware build you install on the management board can either be the same firmware build or a different one.
Installing firmware synchronizes the firmware build and configuration from the management board to the FPCs.
3. Check the synchronization status from the Configuration Sync Monitor or using the `diagnose sys confsync status | grep in_sy` command. The following example FortiGate 6301F output shows that the management board is synchronized with all of the FPCs because each line includes `in_sync=1`.

```
diagnose sys confsync status | grep in_sy
FPC6KFT018901327, Secondary, uptime=3773.96, priority=19, slot_id=1:1, idx=1, flag=0x4, in_sync=1
F6KF31T018900143, Primary, uptime=3837.25, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
FPC6KFT018901372, Secondary, uptime=3774.26, priority=20, slot_id=1:2, idx=1, flag=0x4, in_sync=1
F6KF31T018900143, Primary, uptime=3837.25, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
FPC6KFT018901346, Secondary, uptime=3774.68, priority=21, slot_id=1:3, idx=1, flag=0x4, in_sync=1
F6KF31T018900143, Primary, uptime=3837.25, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
FPC6KFT018901574, Secondary, uptime=3774.19, priority=22, slot_id=1:4, idx=1, flag=0x4, in_sync=1
F6KF31T018900143, Primary, uptime=3837.25, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
FPC6KFT018901345, Secondary, uptime=3773.59, priority=23, slot_id=1:5, idx=1, flag=0x4, in_sync=1
F6KF31T018900143, Primary, uptime=3837.25, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
FPC6KFT018901556, Secondary, uptime=3774.82, priority=24, slot_id=1:6, idx=1, flag=0x4, in_sync=1
F6KF31T018900143, Primary, uptime=3837.25, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
F6KF31T018900143, Primary, uptime=3837.25, priority=1, slot_id=1:0, idx=0, flag=0x0, in_sync=1
FPC6KFT018901327, Secondary, uptime=3773.96, priority=19, slot_id=1:1, idx=1, flag=0x24, in_sync=1
FPC6KFT018901345, Secondary, uptime=3773.59, priority=23, slot_id=1:5, idx=2, flag=0x24, in_sync=1
FPC6KFT018901346, Secondary, uptime=3774.68, priority=21, slot_id=1:3, idx=3, flag=0x24, in_sync=1
FPC6KFT018901372, Secondary, uptime=3774.26, priority=20, slot_id=1:2, idx=4, flag=0x24, in_sync=1
FPC6KFT018901556, Secondary, uptime=3774.82, priority=24, slot_id=1:6, idx=5, flag=0x24, in_sync=1
FPC6KFT018901574, Secondary, uptime=3774.19, priority=22, slot_id=1:4, idx=6, flag=0x24, in_sync=1
```

FPC failover in a standalone FortiGate-6000

A FortiGate-6000 will continue to operate even if one or more FPCs fail. If an FPC stops operating, sessions being processed by that FPC also fail. All new sessions are load balanced to the remaining FPCs. The FortiGate-6000 will continue to operate but with reduced performance because fewer FPCs are operating.

An FPC can fail because of a hardware malfunction, a software problem, or a power supply unit (PSU) failure. The FortiGate-6000 includes three hot-swappable PSUs in a 2+1 redundant configuration. At least two of the PSUs must be operating to provide power to the FortiGate-6000. If only one PSU is operating, only four of the FPCs will continue operating (usually the FPCs in slots 1 to 4). For more information about FPC failure with power loss, see [AC power supply units \(PSUs\)](#).

From the management board GUI dashboard, the Sensor Information dashboard widget displays information about the status of the power supplies. If all power supplies are operating, the widget displays their **Status** as **Normal**.

From the management board CLI, you can use the `execute sensor list` command to verify if the power supplies are operating. The command displays the current status of all FortiGate-6000 sensors including the power supply sensors. Power supply sensor entries should be similar to the following (shown for a FortiGate-6301E). The power supply sensor lines start with `PS{1|2|3}`:

```
65 PS1 VIN          alarm=0 value=122 threshold_status=0
66 PS1 VOUT_12V    alarm=0 value=12.032 threshold_status=0
67 PS1 Temp 1      alarm=0 value=24 threshold_status=0
68 PS1 Temp 2      alarm=0 value=36 threshold_status=0
69 PS1 Fan 1       alarm=0 value=8832 threshold_status=0
70 PS1 Status      alarm=0
71 PS2 VIN          alarm=0 value=122 threshold_status=0
72 PS2 VOUT_12V    alarm=0 value=12.032 threshold_status=0
73 PS2 Temp 1      alarm=0 value=24 threshold_status=0
74 PS2 Temp 2      alarm=0 value=37 threshold_status=0
75 PS2 Fan 1       alarm=0 value=9088 threshold_status=0
76 PS2 Status      alarm=0
77 PS3 VIN          alarm=0 value=122 threshold_status=0
78 PS3 VOUT_12V    alarm=0 value=12.032 threshold_status=0
79 PS3 Temp 1      alarm=0 value=23 threshold_status=0
80 PS3 Temp 2      alarm=0 value=37 threshold_status=0
81 PS3 Fan 1       alarm=0 value=9088 threshold_status=0
82 PS3 Status      alarm=0
```

Any non zero `alarm` or `threshold_status` values indicate a possible problem with that power supply.

If failed FPCs recover, the FortiGate-6000 will attempt to synchronize the configuration of the FPCs with the management board. If there have been few configuration changes, the failed FPCs may be able to become synchronized and operate normally. If there have been many configuration changes or a firmware upgrade, the FortiGate-6000 may not be able to re-synchronize the FPCs without administrator intervention. For example, see [Synchronizing the FPCs with the management board on page 105](#).

You can't replace an FPC that fails because of a hardware failure. Instead, you should RMA the FortiGate-6000.

To show the status of the FPCs, use the `diagnose load-balance status` command. In the command output, if `Status Message` is `Running` the FPC is operating normally. The following example shows the status of FPCs, for a FortiGate-6301F:

```
diagnose load-balance status
=====
MBD SN: F6KF313E17900032
  Primary FPC Blade: slot-2

  Slot 1: FPC6KF3E17900200
    Status:Working  Function:Active
    Link:          Base: Up      Fabric: Up
    Heartbeat: Management: Good  Data: Good
    Status Message:"Running"
  Slot 2: FPC6KF3E17900201
    Status:Working  Function:Active
    Link:          Base: Up      Fabric: Up
    Heartbeat: Management: Good  Data: Good
    Status Message:"Running"
  Slot 3: FPC6KF3E17900207
```

```

Status:Working   Function:Active
Link:           Base: Up           Fabric: Up
Heartbeat: Management: Good      Data: Good
Status Message:"Running"
Slot 4: FPC6KF3E17900219
Status:Working   Function:Active
Link:           Base: Up           Fabric: Up
Heartbeat: Management: Good      Data: Good
Status Message:"Running"
Slot 5: FPC6KF3E17900235
Status:Working   Function:Active
Link:           Base: Up           Fabric: Up
Heartbeat: Management: Good      Data: Good
Status Message:"Running"
Slot 6: FPC6KF3E17900169
Status:Working   Function:Active
Link:           Base: Up           Fabric: Up
Heartbeat: Management: Good      Data: Good
Status Message:"Running"

```

Troubleshooting an FPC failure

This section describes some steps you can use to troubleshoot an FPC failure or to help provide information about the failure to Fortinet Support.

Displaying FPC link and heartbeat status

Start by running the `diagnose load-balance status` command from the management board CLI to check the status of the FPCs. The following output shows the FPC in slot 1 operating normally and a problem with the FPC in slot 2:

```

diagnose load-balance status
=====
MBD SN: F6KF31T018900143
Primary FPC Blade: slot-1

Slot 1: FPC6KFT018901327
Status:Working   Function:Active
Link:           Base: Up           Fabric: Up
Heartbeat: Management: Good      Data: Good
Status Message:"Running"
Slot 2:
Status:Dead      Function:Active
Link:           Base: Up           Fabric: Down
Heartbeat: Management: Failed    Data: Failed
Status Message:"Waiting for management heartbeat."
...

```

If both the base and fabric links are down

If the `diagnose load-balance status` command shows that both the base and fabric links are down, the FPC may be powered off or shut down.

1. From the management board CLI, run the `execute sensor list` command to check the status of the power supplies. Look for the PS1, PS2, and PS3 output lines.

For example, for PS1:

```
...
65 PS1 VIN          alarm=0  value=122  threshold_status=0
66 PS1 VOUT_12V    alarm=0  value=12.032  threshold_status=0
67 PS1 Temp 1      alarm=0  value=26   threshold_status=0
68 PS1 Temp 2      alarm=0  value=38   threshold_status=0
69 PS1 Fan 1       alarm=0  value=8832  threshold_status=0
70 PS1 Status      alarm=0
...
```

If the power supplies are all OK, the output for all of the PS lines should include `Alarm=0` and `Status=0`.

2. If the command output indicates problems with the power supplies, make sure they are all connected to power. If they are connected, there may be a hardware problem. Contact Fortinet Support for assistance.
3. If the power supplies are connected and operating normally, set up two SSH sessions to the management board.
4. From SSH session 1, enter the following command to connect to the FPC console:


```
execute system console-server connect <slot_id>
```
5. Press Enter to see if there is any response.
6. From SSH session 2, use the following commands to power the FPC off and back on:


```
execute load-balance slot power-off <slot_id>
execute load-balance slot power-on <slot_id>
```
7. From SSH session 1, check to see if the FPC starts up normally after running the `power-on` command.
8. If SSH session 1 shows the FPC starting up, when it has fully started, use the `get system status` command to compare the FPC and management board FortiOS versions.

If the versions don't match, see [Updating FPC firmware to match the management board on page 110](#)
9. If the FPC doesn't start up there may be a hardware problem, contact Fortinet Support for assistance.

If only one link is down

If the base or fabric link is up, then check the Heartbeat line of the `diagnose load-balance status` output. The following conditions on the FPC can cause the management heartbeat to fail:

- The FPC did not start up correctly.
- The FPC software may have stopped operating because a process has stopped.
- The FPC may have experienced a kernel panic.
- The FPC may have experienced a daemon or processes panic.

To get more information about the cause:

1. Set up two SSH sessions to the management board.
2. From SSH session 1, enter the following command to connect to the FPC console:


```
execute system console-server connect <slot_id>
```
3. Press Enter to see if there is any response.

4. If there is a response to SSH session 1 and if you can log into the FPC from SSH session 1:
 - a. Dump the crash log by entering:

```
diagnose debug crashlog read
```
 - b. Use the `get system status` command to compare the FPC and management board FortiOS versions.
If the versions don't match, see [Updating FPC firmware to match the management board on page 110](#).
5. If there is no response to SSH session 1, or if you cannot log into the FPC from SSH session 1, switch to SSH session 2.
 - a. From SSH session 2, run the NMI reset command:

```
execute load-balance slot nmi-reset <slot_id>
```
 - b. From SSH session 1, check to see if any messages appear.
 - c. If a kernel panic stack trace is displayed, save it.
The FPC should automatically reboot after displaying the stack trace.
 - d. If nothing happens on SSH session 1, go back to SSH session 2, and run the following commands to power off and power on the FPC:

```
execute load-balance slot power-off <slot_id>
execute load-balance slot power-on <slot_id>
```
 - e. If SSH session 1 shows the FPC starting up, when it has fully started, use the `get system status` command to compare the FPC and management board FortiOS versions.
If the versions don't match, see [Updating FPC firmware to match the management board on page 110](#).
 - f. If the versions match, start an SSH session to log into the FPC, and dump the comlog by entering:

```
diagnose debug comlog read
```


If the comlog was not enabled, it will be empty.
 - g. Also dump the crash log if you haven't been able to do so by entering:

```
diagnose debug crashlog read
```
 - h. Contact Fortinet Support for assistance.
If requested you can provide the comlog and crashlog to help determine the cause of the problem.

Updating FPC firmware to match the management board

Use the steps in [Installing firmware on an individual FPC on page 102](#) to update the firmware running on the FPC to match the firmware running on the management board.

Troubleshooting configuration synchronization issues

After confirming that the management board and the FPC are running the same firmware build, use the following command to determine if configuration synchronization errors remain:

```
diagnose sys confsync status
```

In the command output, `in_sync=1` means the FPC is synchronized and can operate normally, `in_sync=0` means the FPC is not synchronized. If the FPC is up but not synchronized, see [Troubleshooting Tip: FortiGate 7000 Series blade config synchronization issues \(confsync\)](#) for help troubleshooting configuration synchronization issues.

Adjusting global DP3 timers

This section describes the global DP3 timers that you can adjust from the CLI. These timers affect the operation of the Fortigate-6000 DP3 processor.

```
config global
  config system global
    set dp-fragment-timer <timer>
    set dp-pinhole-timer <timer>
    set dp-tcp-normal-timer <timer>
    set dp-udp-idle-timer <timer>
  end
```

`dp-fragment-timer` the time to wait for the next fragment of a fragmented packet. The range is 1 to 65535 seconds. The default is 120 seconds. See [Load balancing TCP, UDP, and ICMP sessions with fragmented packets on page 38](#).

`dp-pinhole-timer` the time to wait to close a pinhole if no more matching traffic that would use the pinhole is received by the DP3 processor. The range is 30 to 120 seconds. The default is 120 seconds.

`dp-tcp-normal-timer` the time to wait before the DP3 processor closes an idle TCP session. The range is 1 to 65535 seconds. The default is 3605 seconds. Some FortiGate-6000 implementations may need to increase this timer if TCP or UDP sessions with NAT enabled are expected to or found to be idle for more than 3605 seconds.

`dp-udp-idle-timer` the time to wait before the DP3 processor closes an idle UDP session. The range is 1 to 86400 seconds. The default is 0, which means the DP3 processor uses the UDP idle timer set by the `udp-idle-timer` option of the `config system global` command.

Changing the FortiGate-6301F and 6501F log disk and RAID configuration

The FortiGate 6001F, FortiGate 6501F, or FortiGate 6301F include two internal 1-TByte log disks. By default the disks are in a RAID-1 configuration. In the RAID-1 configuration you can use the disks for disk logging only. You can use the `execute disk raid` command to disable RAID and use one of the disks for disk logging and the other for other purposes such as disk caching. You can also change the RAID level to RAID-0. Changing the RAID configuration deletes all data from the disks and can disrupt disk logging so a best practice is set the RAID configuration when initially setting up the FortiGate 6001F, FortiGate 6501F, or FortiGate 6301F.

From the CLI you can use the following command to show disk status:

```
execute disk list
```

Use the following command to disable RAID:

```
execute disk raid disable
```

RAID is disabled, the disks are separated and formatted.

Use the following command to change the RAID level to RAID-0:

```
execute disk raid rebuild-level 0
```

The disks are formatted for RAID-0.

Use the following command to rebuild the current RAID partition:

```
execute disk raid rebuild
```

The RAID is rebuilt at the current RAID level.

Use the following command to show RAID status. The following command output shows the disks configured for RAID-1.

```
execute disk raid status
RAID Level: Raid-1
RAID Status: OK
RAID Size: 1000GB

Disk 1: OK Used 953GB
Disk 2: OK Used 953GB
```

Restarting the FortiGate 6000F

To restart the FortiGate 6000F, connect to the management board CLI and enter the `execute reboot` command. After you enter this command, the management board and all of the FPCs restart.

To restart an individual FPC, log in to the CLI of that FPC and run the `execute reboot` command.

Packet sniffing for FPC and management board packets

From the management board CLI, you can access a VDOM and use the `diagnose sniffer packet` command to view or sniff packets processed by the FPCs for this VDOM. To use this command, log into the management board and edit a VDOM. The command output will include packets processed by all of the FPCs in the selected VDOM.

You can also use the `diagnose sniffer packet` command from an individual FPC to view packets processed by that FPC.

From the management board the command syntax is:

```
diagnose sniffer packet <interface> <protocol-filter> <verbose> <count> <timestamp> <frame-
size> <slot>
```

Where:

`<interface>` the name of one or more interfaces on which to sniff for packets. Use `any` to sniff packets for all interfaces.

`<protocol-filter>` a filter to select the protocol for which to view traffic. This can be simple, such as entering `udp` to view UDP traffic or complex to specify a protocol, port, and source and destination interface and so on.

`<verbose>` the amount of detail in the output, and can be:

1. display packet headers only.
2. display packet headers and IP data.
3. display packet headers and Ethernet data (if available).
4. display packet headers and interface names.
5. display packet headers, IP data, and interface names.
6. display packet headers, Ethernet data (if available), and interface names.

`<count>` the number of packets to view. You can enter Ctrl-C to stop the sniffer before the count is reached.

`<timestamp>` the timestamp format, `a` for UTC time and `l` for local time.

<frame-size> the frame size that is printed before truncation. Defaults to the interface MTU.

<slot> the FPC(s) for which to view packets.

- To view packets for one FPC enter the slot number of the FPC.
- To view packets for more than one FPC, enter the slot numbers separated by commas. You can also include a range. For example, to view packets for the FPCs in slots 1, 2, 3, and 6 you can enter 1, 2, 3, 6 or 1-3, 6.
- To view packets for all FPCs, enter `all`.
- If you leave out the <slot> option, you can use the `diagnose sniffer options slot` command to set whether management board packets appear or whether management board and FPC packets appear.

Using the diagnose sniffer options slot command

You can use the `diagnose sniffer options slot` command to control what the `diagnose sniffer packet` command displays if you don't include the <slot> option. The default `diagnose sniffer options slot` setting causes the `diagnose sniffer packet` command to display packets processed by all FPCs and by the management board.

You can use the following command to only display packets processed by the management board:

```
diagnose sniffer options slot current
```

Then the next time you enter the `diagnose sniffer packet` command and leave out the <slot> option, only packets from the management board appear in the command output.

Filtering out internal management traffic

The FortiGate-6000 includes internal interfaces that process internal management and synchronization communication between FortiGate-6000 components. Because this traffic uses internal interfaces, if you specify one or more interface names in the `diagnose sniffer packet` command this traffic is filtered out. However, if you sniff traffic on any interface, internal management traffic can appear in the `diagnose sniffer packet` command output.

The `diagnose sniffer options filter-out-internal-pkts` option if enabled (the default), filters out this internal management traffic. You can disable this option if you want to see the internal management traffic in the `diagnose sniffer packet` output.

Packet sniffing integrated switch fabric (ISF) interfaces

From the management board CLI after editing any VDOM, you can use the following command to sniff traffic on FortiGate 6000F ISF interfaces:

```
diagnose span-sniffer packet <interface> <filter> <verbose> <count> <timestamp> <frame-size>
```

Where:

<interface> the name of one ISF interface on which to sniff for packets. The interface can be `dp` for the ISF interface connected to the DP processor or `sw:port1`, `sw:port2`, `sw:port3`, ..., `sw:port28` for the ISF interfaces connected to traffic interfaces. You must specify an interface name. `any` does not work.

<filter> a filter to select the types of packets for which to view traffic. This can be simple, such as entering `udp` to view UDP traffic or complex to specify a protocol, port, and source and destination interface and so on.

<verbose> the amount of detail in the output, and can be:

1. display packet headers only.
2. display packet headers and IP data.
3. display packet headers and Ethernet data (if available).
4. display packet headers and interface names.
5. display packet headers, IP data, and interface names.
6. display packet headers, Ethernet data (if available), and interface names.

<count> the number of packets to view. You can enter Ctrl-C to stop the sniffer before the count is reached. If you don't include a count packets are displayed continuously until you press Ctrl-C.

<timestamp> the timestamp format, a for UTC time, l for local time, and otherwise to display the time relative to entering the command in the format ss.ms.

<frame-size> the frame size that is printed before truncation. Defaults to the interface MTU.

NMI switch and NMI reset commands

When working with Fortinet Support to troubleshoot problems with your FortiGate-6000 you can use the front panel non-maskable interrupt (NMI) switch to assist with troubleshooting. Pressing this switch causes the software to dump management board registers and backtraces to the console. After the data is dumped, the management board restarts and traffic is temporarily blocked. The management board should restart normally and traffic can resume once the management board is up and running.

You can use the following command to dump registers and backtraces of one or more FPCs to the console. After the data is dumped, the FPC or FPCs reboot. While the FPCs are rebooting, traffic is distributed to the remaining FPCs. The FPCs should restart normally and traffic can resume once they are up and running.

```
execute load-balance slot nmi-reset <slot-number(s)>
```

Where <slot-number(s)> can be one or more FPC slot numbers or slot number ranges with no space and separated by commas. For example:

```
execute load-balance slot nmi-reset 1,3-4
```

Diagnose debug flow trace for FPC and management board activity

The `diagnose debug flow trace` output from the FortiGate-6000 management board CLI now displays debug data for the management board and for all of the FPCs. Each line of output begins with the name of the component that produced the output. For example:

```
diagnose debug enable
[FPC06] id=20085 trace_id=2 func=resolve_ip6_tuple_fast line=4190 msg="vd-vlan:0 received a packet
(proto=6, 3ff5::100:10001->4ff5::13:80) from vlan-port1."
[FPC07] id=20085 trace_id=2 func=resolve_ip6_tuple_fast line=4190 msg="vd-vlan:0 received a packet
(proto=6, 3ff5::100:10000->4ff5::11:80) from vlan-port1."
[FPC06] id=20085 trace_id=2 func=resolve_ip6_tuple line=4307 msg="allocate a new session-000eb730"
[FPC07] id=20085 trace_id=2 func=resolve_ip6_tuple line=4307 msg="allocate a new session-000eb722"
[FPC06] id=20085 trace_id=2 func=vf_ip6_route_input line=1125 msg="find a route: gw-4ff5::13 via vlan-
port2 err 0 flags 01000001"
```

Running FortiGate-6000 `diagnose debug flow trace` commands from an individual FPC CLI shows traffic processed by that FPC only. For example:

```
diagnose debug enable
[FPC02] id=20085 trace_id=2 func=resolve_ip6_tuple_fast line=4190 msg="vd-vlan:0 received a packet
(proto=6, 3ff5::100:10001->4ff5::28:80) from vlan-port1."
[FPC02] id=20085 trace_id=2 func=resolve_ip6_tuple line=4307 msg="allocate a new session-000f00fb"
[FPC02] id=20085 trace_id=2 func=vf_ip6_route_input line=1125 msg="find a route: gw-4ff5::28 via vlan-
port2 err 0 flags 01000001"
[FPC02] id=20085 trace_id=2 func=fw6_forward_handler line=345 msg="Check policy between vlan-port1 ->
vlan-port2"
```

FortiGate-6000 config CLI commands

This chapter describes the following FortiGate-6000 load balancing configuration commands:

- [config load-balance flow-rule](#)
- [config load-balance setting](#)
- [config system console-server](#)

config load-balance flow-rule

Use this command to create flow rules that add exceptions to how matched traffic is processed. You can use flow rules to match a type of traffic and control whether the traffic is forwarded or blocked. And if the traffic is forwarded, you can specify whether to forward the traffic to a specific slot or slots. Unlike firewall policies, load-balance rules are not stateful so for bi-directional traffic, you may need to define two flow rules to match both traffic directions (forward and reverse).

Syntax

```
config load-balance flow-rule
edit <id>
    set status {disable | enable}
    set src-interface <interface-name> [<interface-name>...]
    set vlan <vlan-id>
    set ether-type {any | arp | ip | ipv4 | ipv6}
    set src-addr-ipv4 <ip4-address> <netmask>
    set dst-addr-ipv4 <ip4-address> <netmask>
    set src-addr-ipv6 <ip6-address> <netmask>
    set dst-addr-ipv6 <ip6-address> <netmask>
    set protocol {<protocol-number> | any | icmp | icmpv6 | tcp | udp | igmp | sctp | gre
        | esp | ah | ospf | pim | vrrp}
    set src-l4port <start>[--<end>]
    set dst-l4port <start>[--<end>]
    set icmptype <type>
    set icmpcode <type>
    set tcp-flag {any | syn | fin | rst}
    set action {forward | mirror-ingress | stats | drop}
    set mirror-interface <interface-name>
    set forward-slot {master | all | load-balance | <FPC#>}
    set priority <number>
    set comment <text>
end
```

status {disable | enable}

Enable or disable this flow rule. New flow rules are disabled by default.

src-interface <interface-name> [interface-name>...]

Optionally add the names of one or more front panel interfaces accepting the traffic to be subject to the flow rule. If you don't specify a `src-interface`, the flow rule matches traffic received by any interface.

If you are matching VLAN traffic, select the interface that the VLAN has been added to and use the `vlan` option to specify the VLAN ID of the VLAN interface.

vlan <vlan-id>

If the traffic matching the rule is VLAN traffic, enter the VLAN ID used by the traffic. You must set `src-interface` to the interface that the VLAN interface is added to.

ether-type {any | arp | ip | ipv4 | ipv6}

The type of traffic to be matched by the rule. You can match any traffic (the default) or just match ARP, IP, IPv4 or IPv6 traffic.

{src-addr-ipv4 | dst-addr-ipv4} <ipv4-address> <netmask>

The IPv4 source and destination address of the IPv4 traffic to be matched. The default of `0.0.0.0 0.0.0.0` matches all IPv4 traffic. Available if `ether-type` is set to `ipv4`.

{src-addr-ipv6 | dst-addr-ipv6} <ip-address> <netmask>

The IPv6 source and destination address of the IPv6 traffic to be matched. The default of `::/0` matches all IPv6 traffic. Available if `ether-type` is set to `ipv6`.

protocol {<protocol-number> | any | icmp | icmpv6 | tcp | udp | igmp | sctp | gre | esp | ah | ospf | pim | vrrp}

If `ether-type` is set to `ip`, `ipv4`, or `ipv6`, specify the protocol of the IP, IPv4, or IPv6 traffic to match the rule. The default is `any`. You can specify any protocol number or you can use the following keywords to select common protocols.

Option	Protocol number
icmp	1
icmpv6	58
tcp	6
udp	17
igmp	2
sctp	132
gre	47

Option	Protocol number
esp	50
ah	51
ospf	89
pim	103
vrrp	112

{src-l4port | dst-l4port} <start>[-<end>]

Specify a layer 4 source port range and destination port range. This option appears when `protocol` is set to `tcp` or `udp`. The default range is 0-0, which matches all ports. You don't have to enter a range to match just one port. For example, to set the source port to 80, enter `set src-l4port 80`.

icmp-type <type>

Specify an ICMP type number in the range of 0 to 255. The default is 255. This option appears if `protocol` is set to `icmp`. For information about ICMP type numbers, see [Internet Control Message Protocol \(ICMP\) Parameters](#).

icmp-code <type>

If the ICMP type also includes an ICMP code, you can use this option to add that ICMP code. The ranges is 0 to 255. The default is 255. This option appears if `protocol` is set to `icmp`. For information about ICMP code numbers, see [Internet Control Message Protocol \(ICMP\) Parameters](#).

tcp-flag {any | syn | fin | rst}

Set the TCP session flag to match. The `any` setting (the default) matches all TCP sessions. You can add specific flags to only match specific TCP session types.

action {forward | mirror-ingress | stats | drop}

The action to take with matching sessions. They can be dropped, forwarded to another destination, or you can record statistics about the traffic for later analysis. You can combine two or three settings in one command for example, you can set `action` to both `forward` and `stats` to forward traffic and collect statistics about it. Use `append` to append additional options.

The default action is `forward`, which forwards packets to the specified `forward-slot`.

The `mirror-ingress` option copies (mirrors) all ingress packets that match this flow rule and sends them to the interface specified with the `mirror-interface` option.

set mirror-interface <interface-name>

The name of the interface to send packets matched by this flow-rule to when `action` is set to `mirror-ingress`.

forward-slot {master | all | load-balance | <FPC#>}

The slot that you want to forward the traffic that matches this rule to.

Where:

`master` forwards traffic to the primary FPC.

`all` means forward the traffic to all FPCs.

`load-balance` means forward this traffic to the DP processors that then use the default load balancing configuration to handle this traffic.

`<FPC#>` forward the matching traffic to a specific FPC. For example, FPC3 is the FPC in slot 3.

priority <number>

Set the priority of the flow rule in the range 1 (lowest priority) to 10 (highest priority). Higher priority rules are matched first. You can use the priority to control which rule is matched first if you have overlapping rules.

The default priority is 5.

comment <text>

Optionally add a comment that describes the flow rule.

config load-balance setting

Use this command to set a wide range of load balancing settings.

```
config load-balance setting
  set slbc-mgmt-intf {mgmt1 | mgmt2 | mgmt3}
  set max-miss-heartbeats <heartbeats>
  set max-miss-mgmt-heartbeats <heartbeats>
  set weighted-load-balance {disable | enable}
  set gtp-load-balance {disable | enable}
  set pfc-p-load-balance {disable | enable}
  set sslvpn-load-balance {disable | enable}
  set dp-fragment-session {disable | enable}
  set dp-load-distribution-method {to-master | src-ip | dst-ip | src-dst-ip | src-ip-sport
    | dst-ip-dport | src-dst-ip-sport-dport}
  set sw-load-distribution-method {src-dst-ip | src-dst-ip-sport-dport}
  set dp-icmp-distribution-method {to-master | src-ip | dst-ip | src-dst-ip | derived}
  set dp-session-table-type {intf-vlan-based | vdom-based}
  set nat-source-port {chassis-slots | enabled-slots}
  config workers
    edit <slot>
      set status {disable | enable}
      set weight <weight>
    end
```

slbc-mgmt-intf {mgmt1 | mgmt2 | mgmt3}

Selects the interface used for management connections. The default is `mgmt1`. The IP address of this interface becomes the IP address used to enable management access to individual FPCs using special administration ports as described in [Special management port numbers on page 27](#). To manage individual FPCs, this interface must be connected to a network.



To enable using the special management port numbers to connect to individual FPCs, set `slbc-mgmt-intf` to an interface that is connected to a network, has a valid IP address, and has management or administrative access enabled. To block access to the special management port numbers, you can set `slbc-mgmt-intf` to an interface that is not connected to a network, does not have a valid IP address, or has management or administrative access disabled.

max-miss-heartbeats <heartbeats>

Set the number of missed heartbeats before an FPC is considered to have failed. If a failure occurs, the DP3 processor will no longer load balance sessions to the FPC.

The time between heartbeats is 0.2 seconds. Range is 3 to 300. A value of 3 means 0.6 seconds, 20 (the default) means 4 seconds, and 300 means 60 seconds.

max-miss-mgmt-heartbeats <heartbeats>

Set the number of missed management heartbeats before a FPC is considering to have failed. If a failure occurs, the DP3 processor will no longer load balance sessions to the FPC.

The time between management heartbeats is 1 second. Range is 3 to 300 heartbeats. The default is 10 heartbeats.

weighted-load-balance {disable | enable}

Enable weighted load balancing depending on the slot (or worker) weight. Use `config workers` to set the weight for each slot or worker.

ipsec-load-balance {disable | enable}

Enable or disable IPsec VPN load balancing.

By default IPsec VPN load balancing is enabled and the FortiGate-6000 directs IPsec VPN sessions to the DP3 processors which load balance them among the FPCs.

If IPsec VPN load balancing is enabled, the FortiGate-6000 will drop IPsec VPN sessions traveling between two IPsec tunnels because the two IPsec tunnels may be terminated on different FPCs. If you have traffic entering the FortiGate-6000 from one IPsec VPN tunnel and leaving the FortiGate-6000 out another IPsec VPN tunnel you need to disable IPsec load balancing. Disabling IPsec VPN load balancing enables the default IPsec VPN flow-rules.

gtp-load-balance {disable | enable}

Enable or disable GTP-U load balancing. For more information, see [Enabling GTP load balancing on page 34](#).

pfcp-load-balance {disable | enable}

Enable or disable PFCP user plane load balancing. For more information, see [PFCP load balancing on page 36](#).

set sslvpn-load-balance {disable | enable}

Enable or disable SSL VPN load balancing. For more information, see [SSL VPN load balancing on page 44](#).

dp-fragment-session {disable | enable}

Enable or disable efficient DP3 load balancing of TCP, UDP, and ICMP sessions with fragmented packets. The option is disabled by default.

For more information, see [Load balancing TCP, UDP, and ICMP sessions with fragmented packets on page 38](#).

dp-load-distribution-method {to-master fd| round-robin | src-ip | dst-ip | src-dst-ip | src-ip-sport | dst-ip-dport | src-dst-ip-sport-dport}

Set the method used to load balance sessions among FPCs. Usually you would only need to change the load balancing method if you had specific requirements or you found that the default method wasn't distributing sessions in the manner that you would prefer. The default is `src-dst-ip-sport-dport` which means sessions are identified by their source address and port and destination address and port.

`to-master` directs all session to the primary FPC. This method is for troubleshooting only and should not be used for normal operation. Directing all sessions to the primary FPC will have a negative impact on performance.

`src-ip` sessions are distributed across all FPCs according to their source IP address.

`dst-ip` sessions are statically distributed across all FPCs according to their destination IP address.

`src-dst-ip` sessions are distributed across all FPCs according to their source and destination IP addresses.

`src-ip-sport` sessions are distributed across all FPCs according to their source IP address and source port.

`dst-ip-dport` sessions are distributed across all FPCs according to their destination IP address and destination port.

`src-dst-ip-sport-dport` distribute sessions across all FPCs according to their source and destination IP address, source port, and destination port. This is the default load balance algorithm and represents true session-aware load balancing. Session aware load balancing takes all session information into account when deciding where to send new sessions and where to send additional packets that are part of an already established session.



The `src-ip` and `dst-ip` load balancing methods use layer 3 information (IP addresses) to identify and load balance sessions. All of the other load balancing methods (except for `to-master`) use both layer 3 and layer 4 information (IP addresses and port numbers) to identify a TCP and UDP session. The layer 3 and layer 4 load balancing methods only use layer 3 information for other types of traffic (SCTP, ICMP, and ESP). If GTP load balancing is enabled, Tunnel Endpoint Identifiers (TEIDs) are used to identify GTP sessions.

sw-load-distribution-method {src-dst-ip | src-dst-ip-sport-dport}

Configure the load distribution method used by the Internal Switch Fabric (ISF). The default setting is `src-dst-ip-sport-dport`.

To support load balancing sessions with fragmented packets, set `sw-load-distribution-method` to `src-dst-ip`. For more information, see [Load balancing TCP, UDP, and ICMP sessions with fragmented packets on page 38](#).

dp-icmp-distribution-method {to-master | src-ip | dst-ip | src-dst-ip | derived}

Set the method used to load balance ICMP sessions among FPCs. Usually you would only need to change the load balancing method if you had specific requirements or you found that the default method wasn't distributing sessions in the manner that you would prefer. The default is `to-master`, which means all ICMP sessions are sent to the primary (master) FPC.

`to-master` directs all ICMP session to the primary FPC.

`src-ip` ICMP sessions are distributed across all FPCs according to their source IP address.

`dst-ip` ICMP sessions are statically distributed across all FPCs according to their destination IP address.

`src-dst-ip` ICMP sessions are distributed across all FPCs according to their source and destination IP addresses.

`derived` ICMP sessions are load balanced using the `dp-load-distribution-method` setting. Since port-based ICMP load balancing is not possible, if `dp-load-distribution-method` is set to a load balancing method that includes ports, ICMP load balancing will use the equivalent load balancing method that does not include ports. For example, if `dp-load-distribution-method` is set to the `src-dst-ip-sport-dport` (the default) then ICMP load balancing will use `src-dst-ip` load balancing.

dp-session-table-type {intf-vlan-based | vdom-based}

Change DP processing load balancing mode:

`intf-vlan-based` is the default value and should be used in all cases unless the FortiGate-6000 will support ECMP.

`vdom-based` should only be selected to support ECMP. Enabling VDOM session tables can reduce connections per second (CPS) performance so it should only be enabled if needed to support ECMP. This performance reduction can be more noticeable if the FortiGate-6000 is processing many firewall only sessions. For more information, see [ECMP support on page 90](#).

set nat-source-port {chassis-slots | enabled-slots}

Change SNAT port partitioning behavior. For more information, see [Controlling SNAT port partitioning behavior on page 41](#).

config workers

Set the weight and enable or disable each worker (FPC). Use the `edit` command to specify the slot the FPC is installed in. You can enable or disable each FPC and set a weight for each FPC.

The weight range is 1 to 10. 5 is average (and the default), 1 is -80% of average and 10 is +100% of average. The weights take effect if `weighted-loadbalance` is enabled.

```
config workers
  edit <slot>
    set status enable
    set weight 5
  end
```

config system console-server

Use this command to disable or enable the FortiGate-6000 console server. The console server allows you to use the `execute system console server` command from the management board CLI to access individual FPC consoles in your FortiGate-6000.

Syntax

```
config system console-server
  set status {disable | enable}
  config entries
    edit <slot>
      set slot-id <id>
      set port <port>
    end
```

set status {disable | enable}

Disable or enable the FortiGate-6000 console server. Enabled by default. The `edit <slot>` configuration shows the port number used for each slot. These settings cannot be changed.

FortiGate-6000 execute CLI commands

This chapter describes the FortiGate-6000 execute commands. Many of these commands are only available from the management board CLI.

execute factoryreset-shutdown

You can use this command to reset the configuration of the FortiGate-6000 management board and all of the FPCs before shutting the system down. This command is normally used in preparation for resetting and shutting down a FortiGate 6000F.

execute ha manage <id>

In an HA configuration, use this command to log in to the management board of the secondary FortiGate-6000.

<id> is the ID of the secondary FortiGate-6000. Usually the primary FortiGate-6000 ID is 0 and the secondary ID is 1. You can enter the ? to see the list of IDs that you can connect to.

After you have logged in, you can manage the secondary FortiGate-6000 from the management board or you can use the `execute-load-balance slot manage` command to connect to the different FPCs in the secondary FortiGate-6000.

execute load-balance slot manage {<chassis>.slot | slot}

Log into the CLI of an individual FPC. Use <slot> to specify the FPC slot number.

In an FGCP HA cluster you must also include the <chassis> number which can be 1 or 2.

You will be asked to authenticate to connect to the FPC. Use the `exit` command to end the session and return to the CLI from which you ran the original command.

execute load-balance slot nmi-reset <slot-map>

Perform an NMI reset on selected FPCs. The NMI reset dumps registers and backtraces of one or more FPCs to the console. After the data is dumped, the FPCs reboot. While the FPCs are rebooting, traffic is distributed to the remaining FPCs. The FPCs should restart normally and traffic can resume once they are up and running. You can use the `diagnose sys confsync status` command to verify that the FPCs have started up.

<slot-map> can be one or more FPC slot numbers or slot number ranges with no spaces and separated by commas. For example, to perform an NMI reset of slots 1, 3, 4, and 5, enter

```
execute load-balance slot nmi-reset 1,3-5
```

execute load-balance slot power-off <slot-map>

Power off selected FPCs. This command shuts down the FPC immediately. You can use the `diagnose sys confsync status` command to verify that the management board cannot communicate with the FPCs.

You can use the `execute load-balance slot power-on` command to start up powered off FPCs.

execute load-balance slot power-on <slot-map>

Power on and start up selected FPCs. It may take a few minutes for the FPCs to start up. You can use the `diagnose sys confsync status` command to verify that the FPCs have started up.

execute load-balance slot reboot <slot-map>

Restart selected FPCs. It may take a few minutes for the FPCs to shut down and restart. You can use the `diagnose sys confsync status` command to verify that the FPCs have started up.

execute load-balance slot set-primary-worker <slot>

Force an FPC to always be the primary FPC, <slot> is the FPC slot number.

The change takes place right away and all new primary FPC sessions are sent to the new primary FPC. Sessions that had been processed by the former primary FPC do not switch over, but continue to be processed by the former primary FPC.

This command is most often used for troubleshooting or testing. Since the command does not change the configuration, if the FortiGate-6000 restarts, the usual primary FPC selection process occurs.

execute set-next-reboot rollback

You can use the following command to change the firmware image that the management board and all of the FPCs load the next time the FortiGate-6000 starts up.

```
execute set-next-reboot rollback
```

This command causes each component to select the firmware image stored on its non-active partition the next time the system starts up. The new command replaces the need to log into each component CLI and running the `execute set-next-reboot {primary | secondary}` command.

You can install firmware on the backup partition of the management board or an FPC using the `execute restore secondary-image` command or from the BIOS.

execute system console-server {clearline | connect | showline}

From the management board CLI, the `execute system console server` command provides access to individual FPC consoles in your FortiGate-6000. Console access can be useful for troubleshooting. For example, if an FPC does not boot properly you can use console access to view the state of the FPC and enter commands to fix the problem or restart the FPC.



The `execute system console-server` commands allow access only to FPCs in the FortiGate-6000 that you are logged into. You can't use this command to access FPCs in the other FortiGate-6000 in an HA configuration.

You can use the `config system console-server` command to enable or disable the console server (enabled by default). For more information, see [config system console-server on page 123](#).

execute system console-server clearline <line>

Clear an active console server. You can use this command to stop a console-server session that you have started with the `execute system console-server connect` command. `<line>` is the console server session number. Use the `execute system console-server showline` command to view the active console server sessions.

execute system console-server connect <slot>

Start a console-server connection from the management board CLI to an FPC CLI. `<slot>` is the FPC slot number. Authenticate to log into the console and use CLI commands to view information, make changes, or restart the FPC. When you are done, use **Ctrl-X** to exit from the console back to the management board CLI.

Using **Ctrl-X** may not work if you are accessing the CLI console from the GUI. Instead you may need to log out of the GUI and then log back in.

execute system console-server showline

Show active console-server sessions.

execute upload image {ftp | tftp | usb}

Use this command to upload a firmware image to the FortiGate-6000 internal TFTP server. Once you have uploaded this firmware image, you can install it on an FPC as described in [Installing firmware on an individual FPC on page 102](#).

You can get the firmware image from an external FTP server, an external TFTP server, or from a USB key plugged in the FortiGate-6000 USB port. Use the following syntax:

```
execute upload image ftp <image-file-and-path> <comment> <ftp-server-address> <username>
    <password>
execute upload image tftp <image-file> <comment> <tftp-server-address>
execute upload image usb <image-file-and-path> <comment>
```



www.fortinet.com

Copyright© 2025 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results 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 certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. 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.