FortiGate-7000 - Handbook

Version 5.4.9

FORTINET DOCUMENT LIBRARY
https://docs.fortinet.com

FORTINET VIDEO GUIDE
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/support-and-training/training.html

NSE INSTITUTE
https://training.fortinet.com

FORTIGUARD CENTER
https://fortiguard.com/

END USER LICENSE AGREEMENT
https://www.fortinet.com/doc/legal/EULA.pdf

FEEDBACK
Email: techdoc@fortinet.com

January 22, 2020
FortiGate-7000 5.4.9 Handbook
01-549-396655-20200122

TABLE OF CONTENTS

Change log

Date	Change description
January 22, 2020	Usage-based ECMP load balancing is not supported. See FortiOS features not supported by FortiGate-7000 v5.4.9 on page 89.
December 20, 2019	New information added to FortiGate-7000 IPsec VPN on page 57 and Adding source and destination subnets to IPsec VPN phase 2 configurations on page 57.
October 4, 2018	Minor updates to the 5.4.9 document release.
August 27, 2018	FortiOS 5.4.9 document release. New sections: Setting up management connections on page 36, Setting up HA management connections on page 69, Distributed clustering on page 73, Modifying heartbeat timing on page 75, Viewing interface statistics on page 56, and FortiGate-7000 v5.4.9 special features and limitations on page 87.
August 1, 2018	Minor updates to the v5.4.5 document release.
June 26, 2018	FortiGate-7000 v5.4.5 build 8047 document release. Minor updates.
December 20, 2017	Updated for FortiGate-7000 v5.4.5. New sections include What's new in for FortiGate-7000 v5.4.9 build 8110 on page 9, FortiGate-7000 v5.4.5 special features and limitations on page 92, FortiGate-7000 IPsec VPN on page 57, gtp-load-balance {disable | enable} on page 106, and Operating a FortiGate-7000 on page 51. Also, changes to High availability on page 64. New section Getting started with FortiGate-7000 on page 35.Additional changes and fixes throughout the document.
November 7, 2017	Changes to Getting started with FortiGate-7000 on page 35. Also added a note about the MGMT interface being a static aggregate and not an LACP aggregate.
November 2, 2017	Updated with new information throughout the document including a new HA chapter.
August 30, 2017	Updated with new information throughout the document.
December 1, 2016	Initial Release

Introduction

This document describes what you need to know to get started using a FortiGate-7000 product. Also included are details about CLI commands that are specific to FortiGate-7000 products.

This FortiOS Handbook chapter contains the following sections:

FortiGate-7000 overview provides a quick overview of FortiGate-7000 components.

Getting started with FortiGate-7000 describes how to get started with managing and configuring your FortiGate-7000 product.

Operating a FortiGate-7000 describes some FortiGate-7000 general operating procedures.

FortiGate-7000 IPsec VPN describes FortiGate-7000-specific IPsec VPN considerations.

High availability describes how to set up FortiGate-7000 high availability.

FortiGate-7000 v5.4.9 special features and limitations describes some special limitations of FortiGate-7000 running FortiOS v5.4.9.

FortiGate-7000 v5.4.5 special features and limitations describes some special limitations of FortiGate-7000 running FortiOS v5.4.5.

FortiGate-7000 v5.4.3 special features and limitations describes some special limitations of FortiGate-7000 running FortiOS v5.4.3.

FortiGate-7000 load balancing commands describes FortiGate-7000 load balancing CLI commands.

What's new in for FortiGate-7000 v5.4.9 build 8110

Version 5.4.9 enhancements mainly consisted of adding FortiOS 5.4.9 to the FortiGate-7000 platform. This release also includes bug fixes and improvements and one additional new feature.

Brief logging mode

Brief logging mode, a carrier grade NAT (CGN) feature, removes some fields from log messages to reduce log message size. Smaller log messages reduces disk space usage and also reduces remote logging network bandwidth usage. Brief logging mode is useful if your FortiGate-7000 system generates a large amount of log messages.

Use the following command to enable brief logging mode for all logging:

```
config log setting
   set brief-traffic-format enable
end
```

What's new in for FortiGate-7000 v5.4.5 build 8047

The following new features have been added to FortiGate-7000 v5.4.5.

The maximum value for user groups has been increased to 5,000 (460857)

On a FortiGate-7000, you can now configure up to 5,000 user groups.

Log field extension policy-name and meta-field (461783 455441)

An option to include the policy name field has been added to traffic logs (log-policy-name). An option to add a meta-field tag to all logs has also been added (custom-field and custom-log-fields; see below). This meta-field could be used to identify the FortiGate sending the logs, for example:

```
config log setting
   set log-policy-name enable
end
config log custom-field
   edit "cust-field"
      set name "MyFortiGate"
      set value "111"
   end
config log setting
   set custom-log-fields "cust-field"
end
```

M1 and M2 interfaces can use different VLANs for heartbeat traffic (408386)

The M1 and M2 interfaces can be configured to use different VLANs for HA heartbeat traffic. Normally you would separate the M1 and M2 traffic. In that case you don't have to change their VLAN IDs. But if the M1 and M2 interfaces are connected to the same switch and you can't separate their traffic, or if you can't use the default VLAN IDs you can set a different VLAN ID for each interface.

Use the following command set the M1 and M2 interfaces to use different VLANs:

```
config system ha
   set hbdev M1/M2
   set hbdev-vlan-id 991
   set hbdev-second-vlan-id 992
end
```

For this configuration to work the `hbdev-vlan-id` has to be changed. You cannot use the default value of 999.

GTP load balancing

GTP load balancing is supported for FortiGate-7000 configurations licensed for FortiOS Carrier. You can use the following command to enable GTP load balancing. This command is only available after you have licensed the FortiGate-7000 for FortiOS Carrier.

```
config load-balance setting
```

```
   set gtp-load-balance enable
end
```

FSSO user authentication synchronization

FSSO user authentication is synchronized to all FPMs. Users authenticated through FSSO no longer have to re-authenticate when load balancing distributes their session to a different FPM.

HA link failure threshold changes (422264)

The link failure threshold is now determined based on the all FIMs in a chassis. This means that the chassis with the fewest active links will become the backup chassis.

FortiGate-7000s running FortiOS v5.4.5 can be configured as dialup IPsec VPN servers

The following shows how to setup a dialup IPsec VPN configuration where the FortiGate-7000 running v5.4.5 acts as a dialup IPsec VPN server.

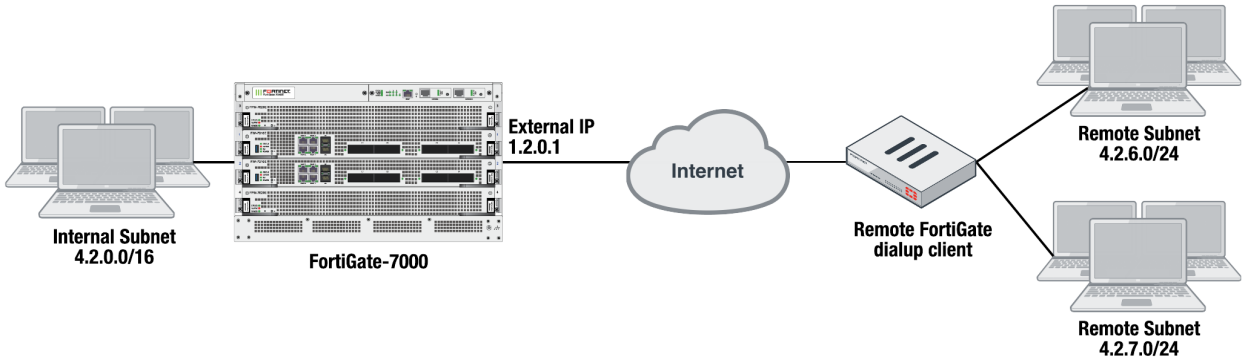

Configure the phase1, set type to dynamic.

```
config vpn ipsec phase1-interface
   edit dialup-server
      set type dynamic
      set interface "v0020"
      set peertype any
      set psksecret < password>
   end
```

Configure the phase 2, to support dialup IPsec VPN, set the destination subnet to 0.0.0.0 0.0.0.0.

```
config vpn ipsec phase2-interface
   edit dialup-server
      set phase1name dialup-server
      set src-subnet 4.2.0.0 255.255.0.0
      set dst-subnet 0.0.0.0 0.0.0.0
   end
```

To configure the remote FortiGate as a dialup IPsec VPN client

The dialup IPsec VPN client should advertise its local subnet(s) using the phase 2 src-subnet option.

If there are multiple local subnets create a phase 2 for each one. Each phase 2 only advertises one local subnet to the dialup IPsec VPN server. If more than one local subnet is added to the phase 2, only the first one is advertised to the server.

Dialup client configuration:

```
config vpn ipsec phase1-interface
   edit "to-fgt7k"
      set interface "v0020"
      set peertype any
      set remote-gw 1.2.0.1
      set psksecret <password>
   end
config vpn ipsec phase2-interface
   edit "to-fgt7k"
      set phase1name "to-fgt7k"
      set src-subnet 4.2.6.0 255.255.255.0
      set dst-subnet 4.2.0.0 255.255.0.0
   next
   edit "to-fgt7k-2"
      set phase1name "to-fgt7k"
      set src-subnet 4.2.7.0 255.255.255.0
      set dst-subnet 4.2.0.0 255.255.0.0
   end
```

FortiGate-7000 overview

A FortiGate-7000 product consists of a FortiGate-7000 series chassis (for example, the FortiGate-7040E) with FortiGate-7000 modules installed in the chassis slots. A FortiGate-7040E chassis comes with two interface modules (FIM) to be installed in slots 1 and 2 to provide network connections and session-aware load balancing to two processor modules (FPM) to be installed in slots 3 and 4.

FortiGate-7000 products are sold and licensed as packages that include the chassis as well as the modules to be included in the chassis. When you receive your FortiGate-7000 series product the chassis has to be installed in a rack and the modules installed in the chassis. Interface modules always go in slots 1 and 2 and processor modules in slots 3 and up.

If your FortiGate-7000 product includes two different interfaces modules, for optimal configuration you should install the module with the lower model number in slot 1 and the module with the higher model number in slot 2. For example, if your chassis includes a FIM-7901E and a FIM-7904E, install the FIM-7901E in chassis slot 1 and the FIM-7904E in chassis slot 2. This applies to any combination of two different interface modules.

As an administrator, when you browse to the FortiGate-7000 management IP address you log into the interface module in slot 1 (the primary or master interface module or FIM) to view the status of the FortiGate-7000 and make configuration changes. The FortiOS firmware running on each module has the same configuration and when you make configuration changes to the primary interface module, the configuration changes are synchronized to all modules.

The same FortiOS firmware build runs on each module in the chassis. You can upgrade FortiGate-7000 firmware by logging into the primary interface module and performing a firmware upgrade as you would for any FortiGate. During the upgrade process the firmware of all of the modules in the chassis upgrades in one step. Firmware upgrades should be done during a quiet time because traffic will briefly be interrupted during the upgrade process.

Licenses, device registration, and support

A FortiGate-7000 product is made up of a FortiGate-7000 series chassis, one or two FIM interface modules and two to four FPM processor modules. The entire package is licensed and configured as a single product under the FortiGate-7000 chassis serial number. When you receive a new FortiGate-7000 product you register it on https://support.fortinet.com using the chassis serial number. Use the chassis serial number when requesting support from Fortinet for the product.

All Fortinet licensing, including FortiCare Support, IPS, AntiVirus, Web Filtering, Mobile Malware, FortiClient, FortiCloud, and additional virtual domains (VDOM) is for the entire FortiGate-7000 product and not for individual components.

If an individual component, such as a single interface or processor fails you can RMA and replace just that component.

FortiGate-7060E

The FortiGate-7060E is a 8U 19-inch rackmount 6-slot chassis with a 80Gbps fabric and 1Gbps base backplane designed by Fortinet. The fabric backplane provides network data communication and the base backplane provides management and synch communication among the chassis slots.

FortiGate-7060E front panel

The chassis is managed by two redundant management modules. Each module includes an Ethernet connection as well as two switchable console ports that provide console connections to the modules in the chassis slots. The active management module controls chassis cooling and power management and provides an interface for managing the modules installed in the chassis.

FortiGate-7060E front panel, (example module configuration)

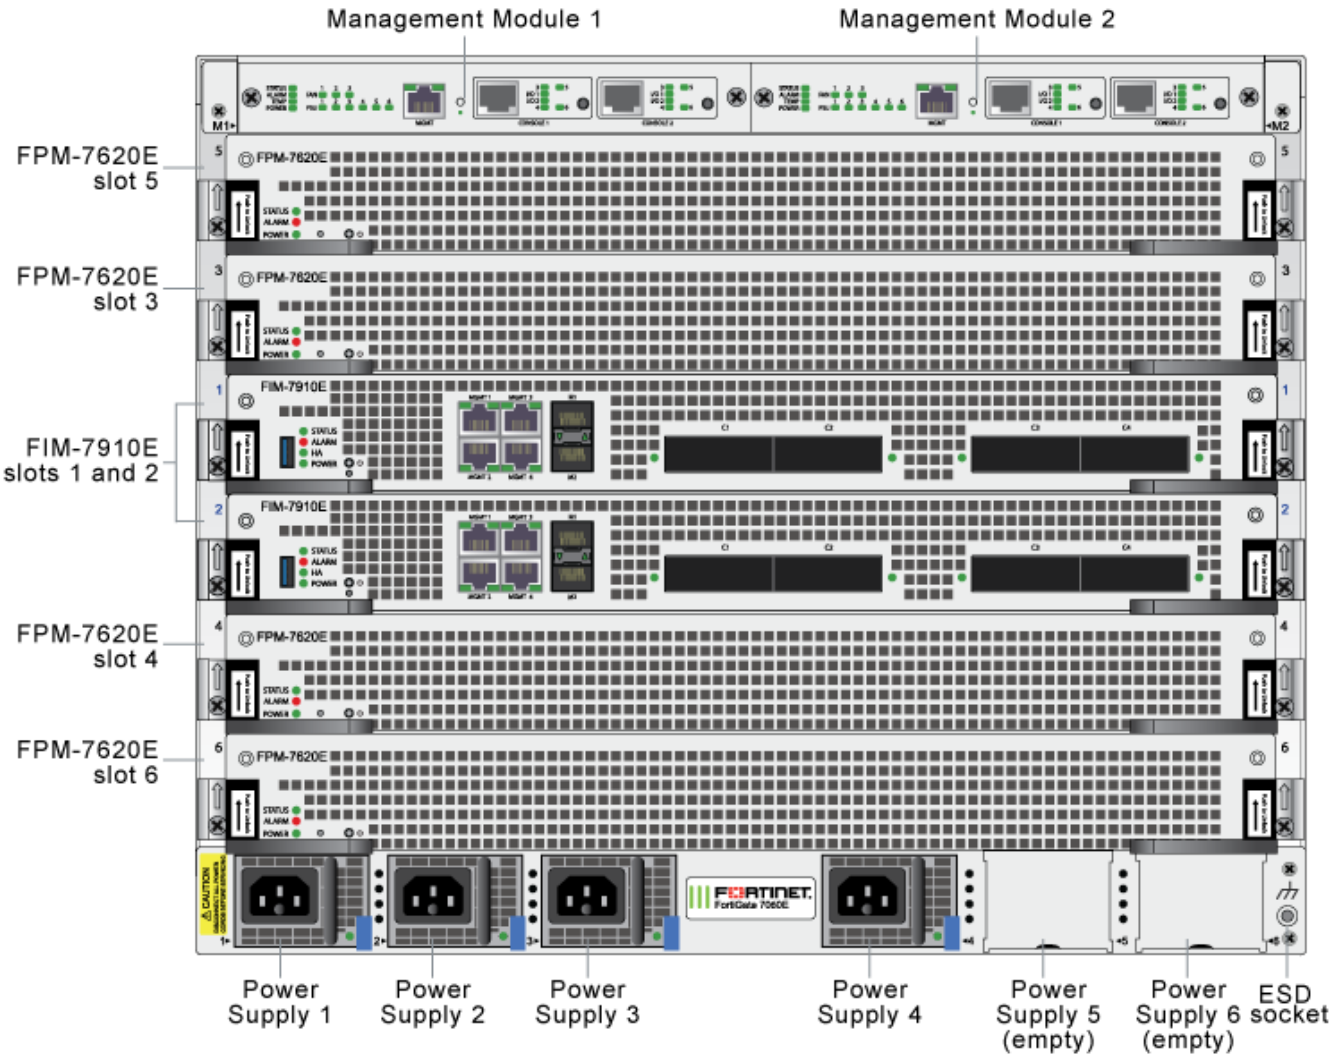

Power is provided to the chassis using four hot swappable 3+1 redundant 100-240 VAC, 50-60 Hz power supply units (PSUs). You can also optionally add up to six PSUs to provide 3+3 redundancy. The FortiGate-7060E can also be equipped with DC PSUs allowing you to connect the chassis to -48V DC power

The standard configuration of the FortiGate-7060E includes two FIM (interface) modules in chassis slots 1 and 2 and up to four FPM (processing) modules in chassis slots 3 to 6.

FortiGate-7060E schematic

The FortiGate-7060E chassis schematic below shows the communication channels between chassis components including the management modules (MGMT), the FIMs (called FIM1 and FIM2) and the FPMs (FPM3, FPM4, FPM5, and FPM6).

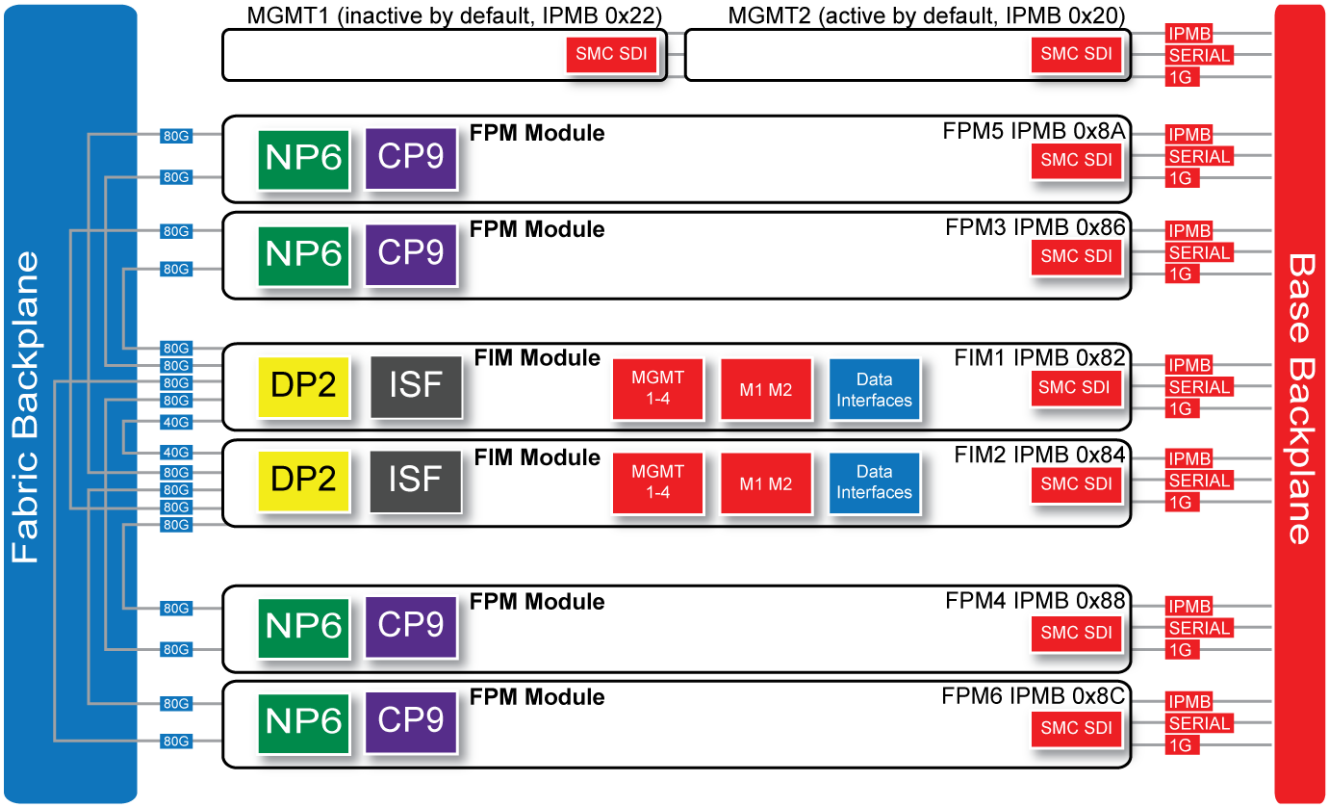

By default, MGMT2 is the active management module and MGMT1 is inactive. The active management module always has the Intelligent Platform Management Bus (IPMB) address 0x20 and the inactive management module always has the IPMB address 0x22.

The active management module communicates with all modules in the chassis over the base backplane. Each module, including the management modules has a Shelf Management Controller (SMC). These SMCs support IPMB communication between the active management module and the FIM and FPMs for storing and sharing sensor data that the management module uses to control chassis cooling and power distribution. The base backplane also supports serial communications to allow console access from the management module to all modules, and 1Gbps Ethernet communication for management and heartbeat communication between modules.

FIM1 and FIM2 (IPMB addresses 0x82 and 0x84) are the FIMs in slots 1 and 2. The interfaces of these modules connect the chassis to data networks and can be used for Ethernet management access to chassis components. The FIMs include DP2 processors that distribute sessions over the Integrated Switch Fabric (ISF) to the NP6 processors in the FPMs. Data sessions are communicated to the FPMs over the 80Gbps chassis fabric backplane.

FPM03, FPM04, FPM05, and FPM06 (IPMB addresses 0x86, 0x88, 0x8A, and 0x8C) are the FPM processor modules in slots 3 to 6. These worker modules process sessions distributed to them by the FIMs. FPMs include NP6 processors to offload sessions from the FPM CPU and CP9 processors that accelerate content processing.

FortiGate-7040E

The FortiGate-7040E is a 6U 19-inch rackmount 4-slot chassis with a 80Gbps fabric and 1Gbps base backplane designed by Fortinet. The fabric backplane provides network data communication and the base backplane provides management and synch communication among the chassis slots.

FortiGate-7040E front panel

The FortiGate-7040E chassis is managed by a single management module that includes an Ethernet connection as well as two switchable console ports that provide console connections to the modules in the chassis slots. The management module controls chassis cooling and power management and provides an interface for managing the modules installed in the chassis. The standard configuration of the FortiGate-7040E includes two FIM (interface) modules in chassis slots 1 and 2 and two FPM (processing) modules in chassis slots 3 and 4.

FortiGate-7040E front panel

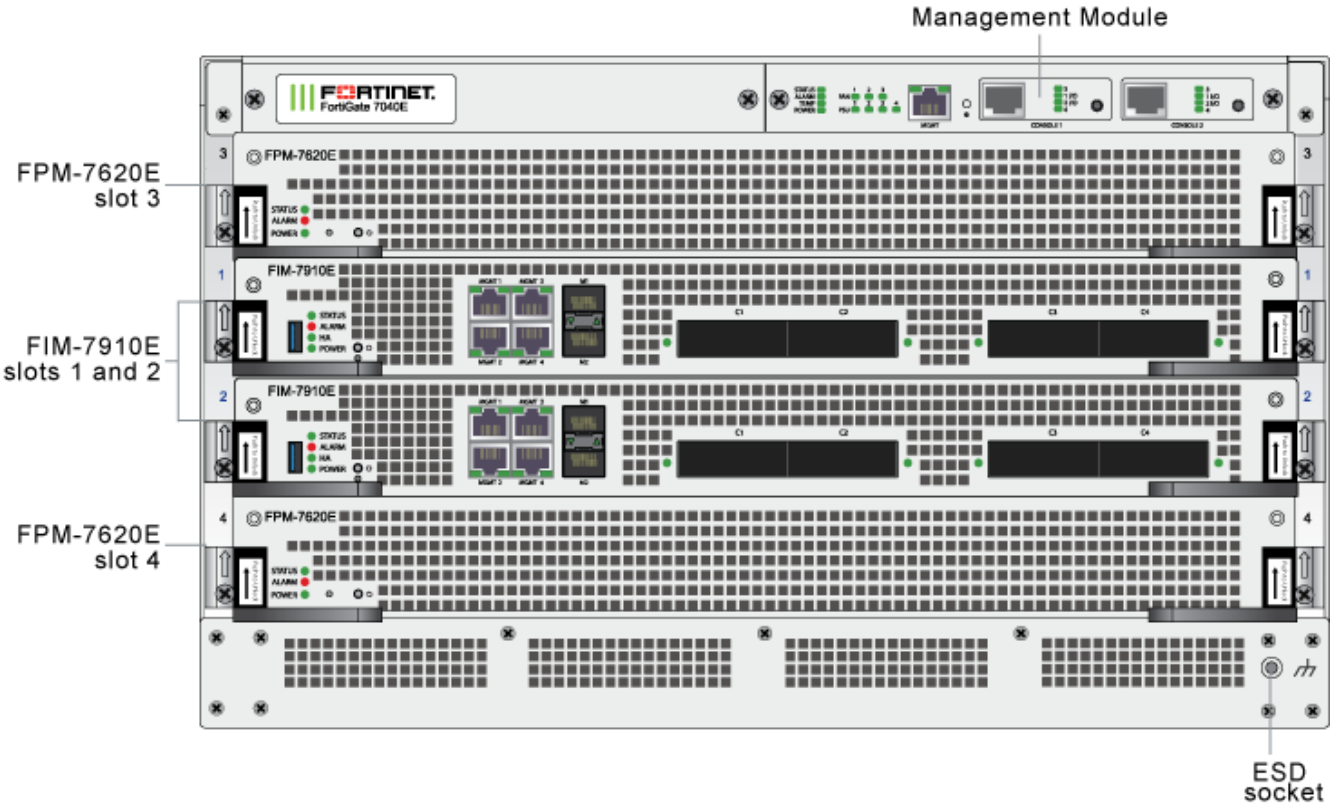

FortiGate-7040E schematic

The FortiGate-7040E chassis schematic below shows the communication channels between chassis components including the management module (MGMT), the FIMs (called FIM1 and FIM2) and the FPMs (FPM3 and FPM4).

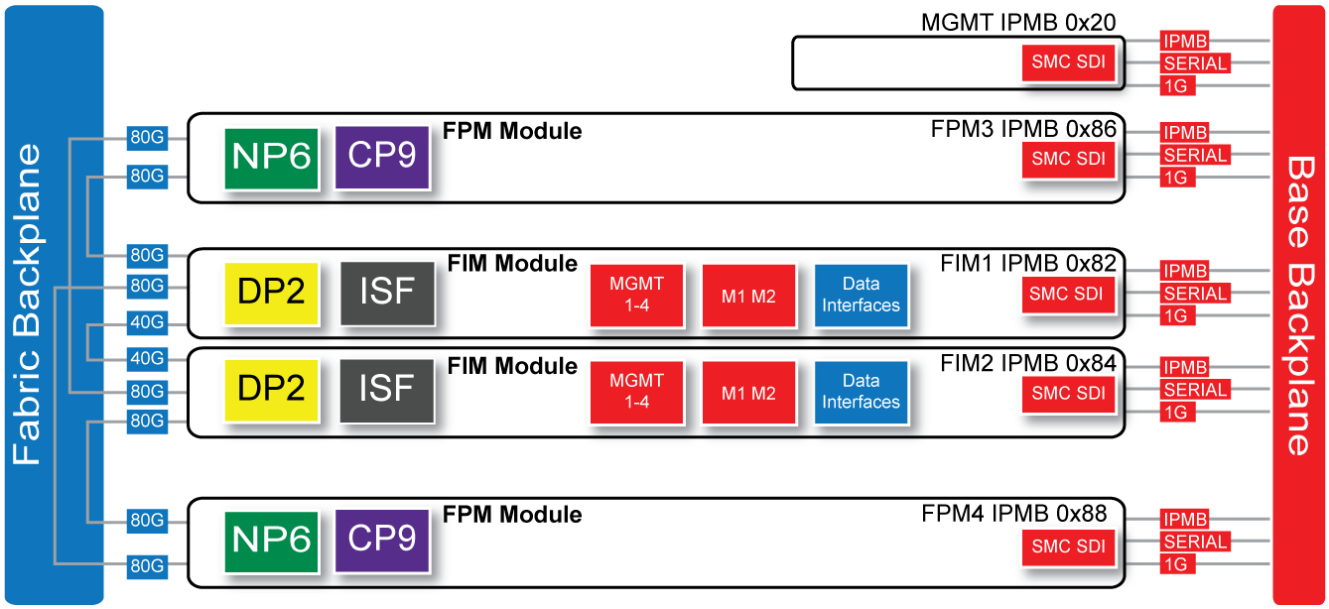

The management module (MGMT, with Intelligent Platform Management Bus (IPMB) address 0x20) communicates with all modules in the chassis over the base backplane. Each module, including the management module, includes a Shelf Management Controller (SMC). These SMCs support IPMB communication between the management module and the FIM and FPMs for storing and sharing sensor data that the management module uses to control chassis cooling and power distribution. The base backplane also supports serial communications to allow console access from the management module to all modules, and 1Gbps Ethernet communication for management and heartbeat communication between modules.

FIM1 and FIM2 (IPMB addresses 0x82 and 0x84) are the FIMs in slots 1 and 2. The interfaces of these modules connect the chassis to data networks and can be used for Ethernet management access to chassis components. The FIMs include DP2 processors that distribute sessions over the Integrated Switch Fabric (ISF) to the NP6 processors in the FPMs. Data sessions are communicated to the FPMs over the 80Gbps chassis fabric backplane.

FPM3 and FPM4 (IPMB addresses 0x86 and 0x88) are the FPM processor modules in slots 3 and 4. These worker modules process sessions distributed to them by the FIMs. FPMs include NP6 processors to offload sessions from the FPM CPU and CP9 processors that accelerate content processing.

FortiGate-7030E

The FortiGate-7030E is a 6U 19-inch rackmount 3-slot chassis with a 80Gbps fabric and 1Gbps base backplane designed by Fortinet. The fabric backplane provides network data communication and the base backplane provides management and synch communication among the chassis slots.

FortiGate-7030E front panel

The FortiGate-7030E chassis is managed by a single management module that includes an Ethernet connection as well as two switchable console ports that provide console connections to the modules in the chassis slots. The management module controls chassis cooling and power management and provides an interface for managing the modules installed in the chassis. The standard configuration of the FortiGate-7030E includes one FIM (interface) module in chassis slot 1 and two FPM (processing) modules in chassis slots 3 and 4. The front panel also includes a sealed blank panel. Breaking the seal or removing the panel voids your FortiGate-7030E warranty.

FortiGate-7030E front panel (example module configuration)

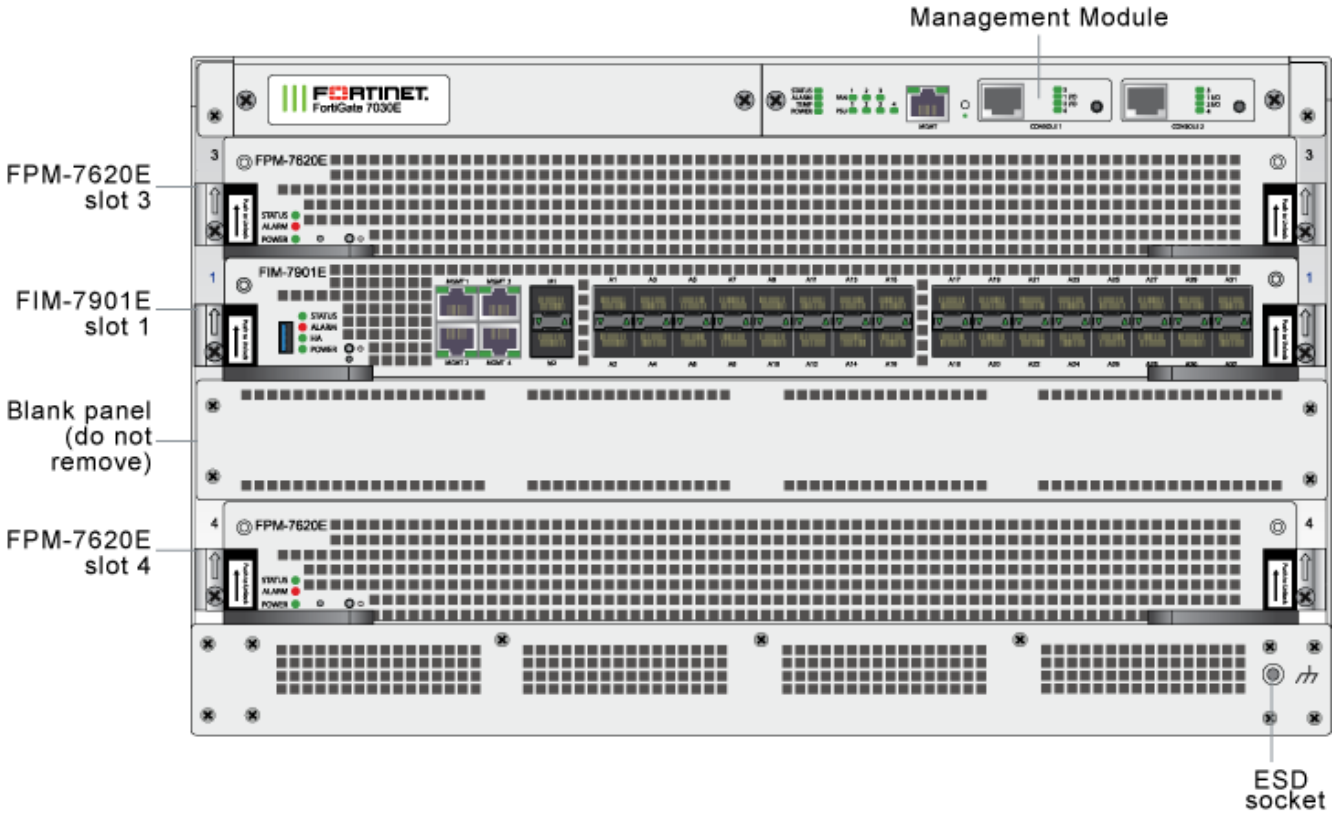

FortiGate-7030E schematic

The FortiGate-7030E chassis schematic below shows the communication channels between chassis components including the management module (MGMT), the FIM (called FIM1) and the FPMs (FPM3 and FPM4).

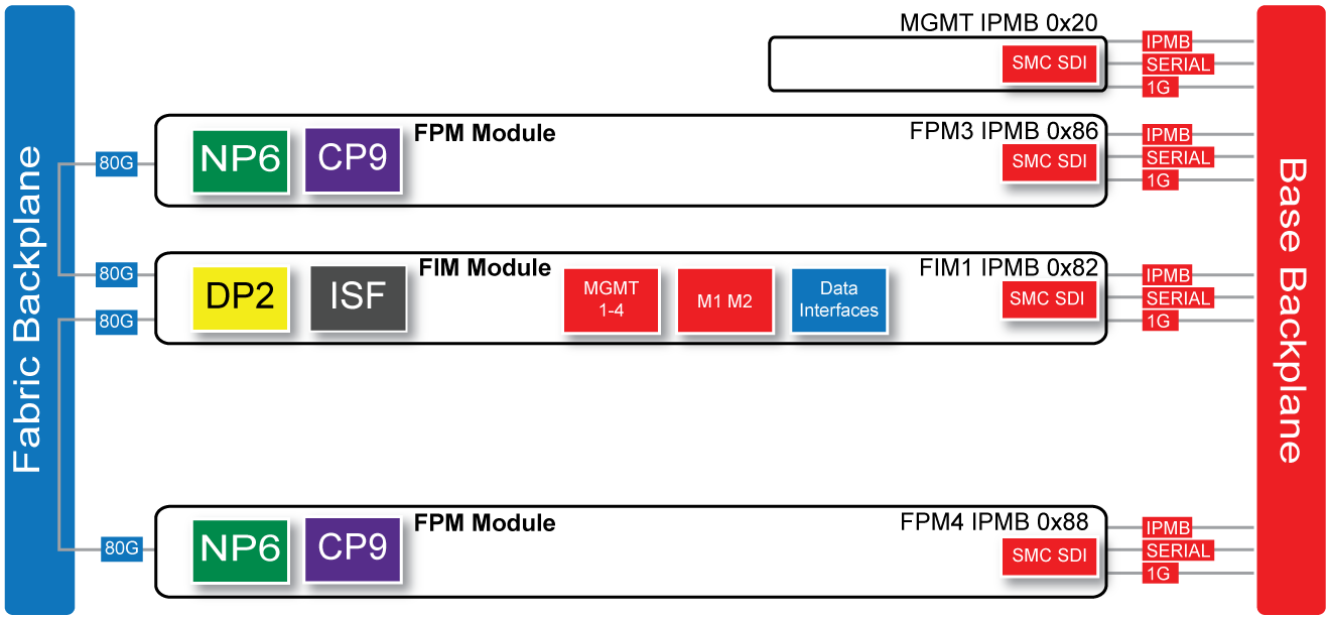

The management module (MGMT, with Intelligent Platform Management Bus (IPMB) address 0x20) communicates with all modules in the chassis over the base backplane. Each module, including the management module includes a Shelf Management Controller (SMC). These SMCs support IPMB communication between the management module and the FIM and FPMs for storing and sharing sensor data that the management module uses to control chassis cooling and power distribution. The base backplane also supports serial communications to allow console access from the management module to all modules, and 1Gbps Ethernet communication for management and heartbeat communication between modules.

FIM1 (IPMB address 0x82) is the FIM in slot 1. The interfaces of this module connect the chassis to data networks and can be used for Ethernet management access to chassis components. The FIM includes DP2 processors that distribute sessions over the Integrated Switch Fabric (ISF) to the NP6 processors in the FPMs. Data sessions are communicated to the FPMs over the 80Gbps chassis fabric backplane.

FPM3 and FPM4 (IPMB addresses 0x86 and 0x88) are the FPM processor modules in slots 3 and 4. These worker modules process sessions distributed to them by the FIM. FPMs include NP6 processors to offload sessions from the FPM CPU and CP9 processors that accelerate content processing.

FIM-7901E interface module

The FIM-7901E interface module is a hot swappable module that provides data, management and session sync/heartbeat interfaces, base backplane switching and fabric backplane session-aware load balancing for a FortiGate-7000 chassis. The FIM-7901E includes an integrated switch fabric and DP2 processors to load balance millions of data sessions over the chassis fabric backplane to FPM processor modules.

The FIM-7901E can be installed in any FortiGate-7000 series chassis in hub/switch slots 1 and 2. The FIM-7901E provides thirty-two 10GigE small form-factor pluggable plus (SPF+) interfaces for a FortiGate-7000 chassis.

You can also install FIM-7901Es in a second chassis and operate the chassis in HA mode with another set of processor modules to provide chassis failover protection.

FIM-7901E front panel

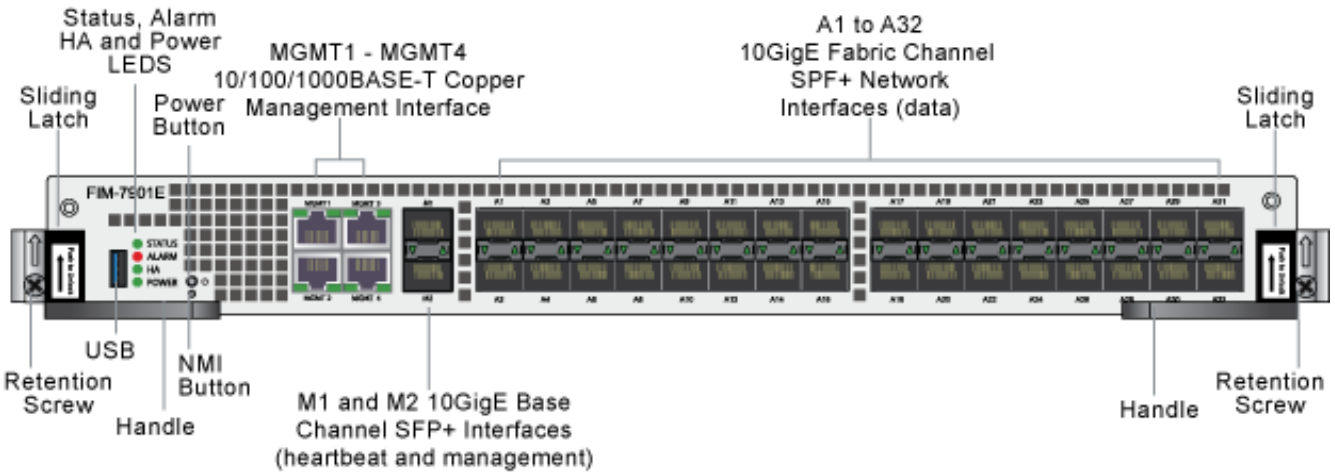

The FIM-7901E includes the following hardware features:

- Thirty-two front panel 10GigE SFP+ fabric channel interfaces (A1 to A32). These interfaces are connected to 10Gbps networks to distribute sessions to the FPM processor modules installed in chassis slots 3 and up. These interfaces can also be configured to operate as Gigabit Ethernet interfaces using SFP transceivers. These interfaces also support creating link aggregation groups (LAGs) that can include interfaces from both FIM-7901Es.
- Two front panel 10GigE SFP+ interfaces (M1 and M2) that connect to the base backplane channel. These interfaces are used for heartbeat, session sync, and management communication between FIM-7901Es in different chassis. These interfaces can also be configured to operate as Gigabit Ethernet interfaces using SFP transceivers, but should not normally be changed. If you use switches to connect these interfaces, the switch ports should be able to accept packets with a maximum frame size of at least 1526. The M1 and M2 interfaces need to be on different broadcast domains. If M1 and M2 are connected to the same switch, Q-in-Q must be enabled on the switch.
- Four 10/100/1000BASE-T out of band management Ethernet interfaces (MGMT1 to MGMT4).
- One 80Gbps fabric backplane channel for traffic distribution with each FPM installed in the same chassis as the FIM-7901E.
- One 1Gbps base backplane channel for base backplane with each FPM installed in the same chassis as the FIM-7901E.
- One 40Gbps fabric backplane channel for fabric backplane communication with the other FIM-7901E in the chassis.
- One 1Gbps base backplane channel for base backplane communication with the other FIM-7901E in the chassis.
- On-board DP2 processors and an integrated switch fabric to provide high-capacity session-aware load balancing.
- One front panel USB port.
- Power button.
- NMI switch (for troubleshooting as recommended by Fortinet Support).
- Mounting hardware.
- LED status indicators.

FIM-7901E schematic

The FIM-7901E includes an integrated switch fabric (ISF) that connects the FIM front panel interfaces to the DP2 session-aware load balancers and to the chassis fabric and base backplanes. The ISF also allows the DP2 processors to

distribute sessions among all NP6 processors on the FortiGate Processor Modules (FPM) in the FortiGate-7000 chassis.

FIM-7901E schematic

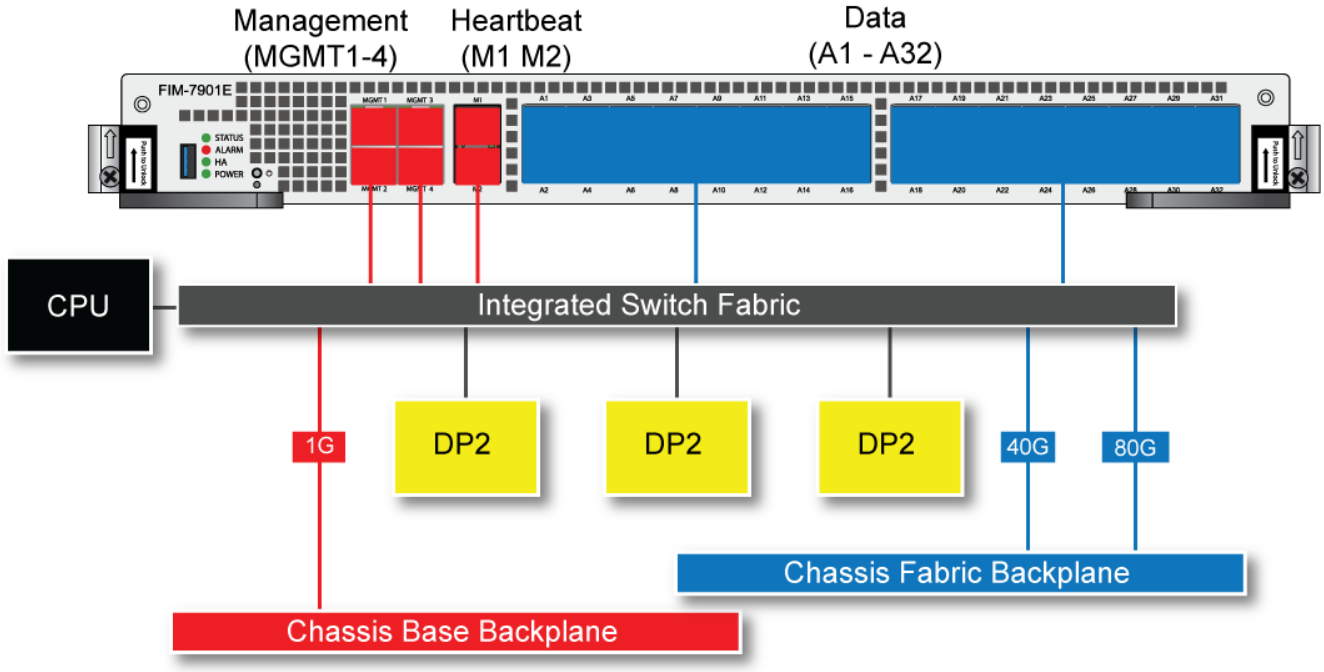

FIM-7904E interface module

The FIM-7904E interface module is a hot swappable module that provides data, management and session sync/heartbeat interfaces, base backplane switching and fabric backplane session-aware load balancing for a FortiGate-7000 series chassis. The FIM-7904E includes an integrated switch fabric and DP2 processors to load balance millions of data sessions over the chassis fabric backplane to FPM processor modules.

The FIM-7904E can be installed in any FortiGate-7000 series chassis in hub/switch slots 1 and 2. The FIM-7904E provides four Quad Small Form-factor Pluggable plus (QSFP+) interfaces for a FortiGate-7000 chassis. Using a 40GBASE-SR10 multimode QSFP+ transceiver, each QSFP+ interface can also be split into four 10GBASE-SR interfaces.

You can also install FIM-7904Es in a second chassis and operate the chassis in HA mode with another set of processor modules to provide chassis failover protection.

FIM-7904E front panel

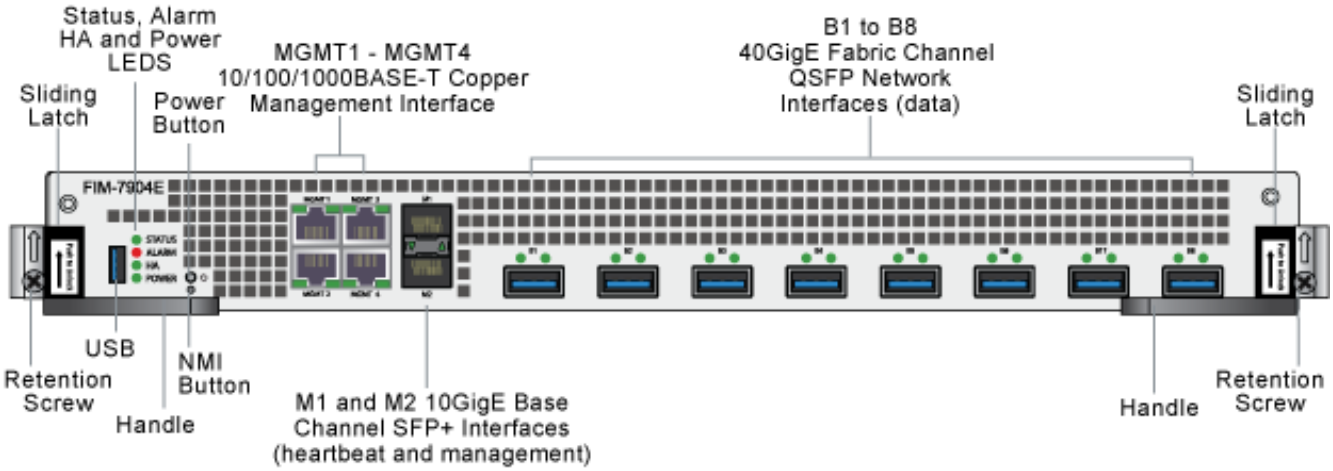

The FIM-7904E includes the following hardware features:

- Eight front panel 40GigE QSFP+ fabric channel interfaces (B1 to B8). These interfaces are connected to 40Gbps networks to distribute sessions to the FPM processor modules installed in chassis slots 3 and up. Using 40GBASE-SR10 multimode QSFP+ transceivers, each QSFP+ interface can also be split into four 10GBASE-SR interfaces. These interfaces also support creating link aggregation groups (LAGs) that can include interfaces from both FIM-7904Es.
- Two front panel 10GigE SFP+ interfaces (M1 and M2) that connect to the base backplane channel. These interfaces are used for heartbeat, session sync, and management communication between FIM-7904Es in different chassis. These interfaces can also be configured to operate as Gigabit Ethernet interfaces using SFP transceivers, but should not normally be changed. If you use switches to connect these interfaces, the switch ports should be able to accept packets with a maximum frame size of at least 1526. The M1 and M2 interfaces need to be on different broadcast domains. If M1 and M2 are connected to the same switch, Q-in-Q must be enabled on the switch.
- Four 10/100/10000BASE-T out of band management Ethernet interfaces (MGMT1 to MGMT4).
- One 80Gbps fabric backplane channel for traffic distribution with each FPM installed in the same chassis as the FIM-7904E.
- One 1Gbps base backplane channel for base backplane with each FPM installed in the same chassis as the FIM-7904E.
- One 40Gbps fabric backplane channel for fabric backplane communication with the other FIM-7904E in the chassis.
- One 1Gbps base backplane channel for base backplane communication with the other FIM-7904E in the chassis.
- On-board DP2 processors and an integrated switch fabric to provide high-capacity session-aware load balancing.
- One front panel USB port.
- Power button.
- NMI switch (for troubleshooting as recommended by Fortinet Support).
- Mounting hardware.
- LED status indicators.

Splitting the FIM-7904E B1 to B8 interfaces

Each 40GE interface (B1 to B8) on the FIM-7904Es in slot 1 and slot 2 of a FortiGate-7000 system can be split into 4x10GBE interfaces. You split these interfaces after the FIM-7904Es are installed in your FortiGate-7000 system and the system us up and running. You can split the interfaces of the FIM-7904Es in slot 1 and slot 2 at the same time by entering a single CLI command. Splitting the interfaces requires a system reboot so Fortinet recommends that you split multiple interfaces at the same time according to your requirements to avoid traffic disruption.

For example, to split the B1 interface of the FIM-7904E in slot 1 (this interface is named 1-B1) and the B1 and B4 interfaces of the FIM-7904E in slot 2 (these interfaces are named 2-B1 and 2-B4) connect to the CLI of your FortiGate-7000 system using the management IP and enter the following command:

```
config system global
   set split-port 1-B1 2-B1 2-B4
end
```

After you enter the command, the FortiGate-7000 reboots and when it comes up:

- The 1-B1 interface will no longer be available. Instead the 1-B1/1, 1-B1/2, 1-B1/3, and 1-B1/4 interfaces will be available.
- The 2-B1 interface will no longer be available. Instead the 2-B1/1, 2-B1/2, 2-B1/3, and 2-B1/4 interfaces will be available.
- The 2-B4 interface will no longer be available. Instead the 2-B4/1, 2-B4/2, 2-B4/3, and 2-B4/4 interfaces will be available.

You can now connect breakout cables to these interfaces and configure traffic between them just like any other FortiGate interface.

FIM-7904E hardware schematic

The FIM-7904E includes an integrated switch fabric (ISF) that connects the front panel interfaces to the DP2 session-aware load balancers and to the chassis backplanes. The ISF also allows the DP2 processors to distribute sessions among all NP6 processors on the FortiGate Processor Modules (FPM) in the FortiGate-7000 chassis.

FIM-7904E hardware architecture

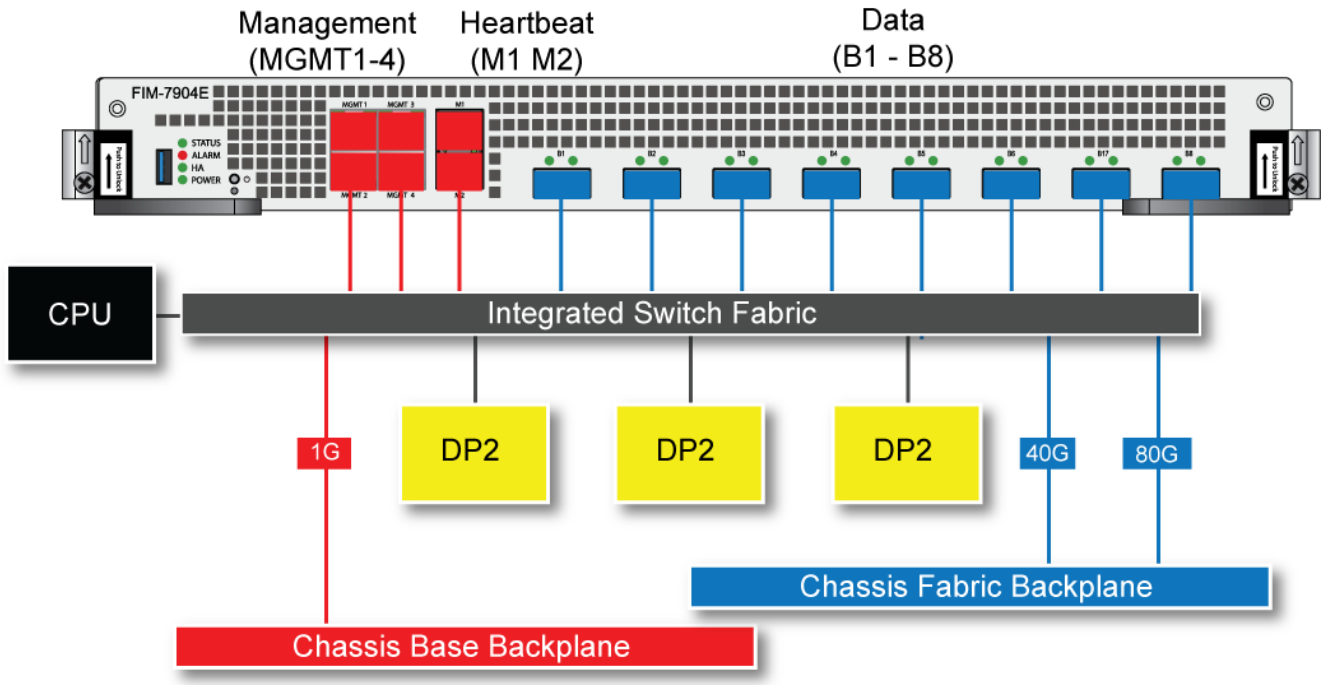

FIM-7910E interface module

The FIM-7910E interface module is a hot swappable module that provides data, management and session sync/heartbeat interfaces, base backplane switching and fabric backplane session-aware load balancing for a FortiGate-7000 series chassis. The FIM-7910E includes an integrated switch fabric and DP2 processors to load balance millions of data sessions over the chassis fabric backplane to FPM processor modules.

The FIM-7910E can be installed in any FortiGate-7000 series chassis in hub/switch slots 1 and 2. The FIM-7910E provides four C form-factor pluggable 2 (CFP2) interfaces for a FortiGate-7000 chassis. Using a 100GBASE-SR10 multimode CFP2 transceiver, each CFP2 interface can also be split into ten 10GBASE-SR interfaces.

FIM-7910E front panel

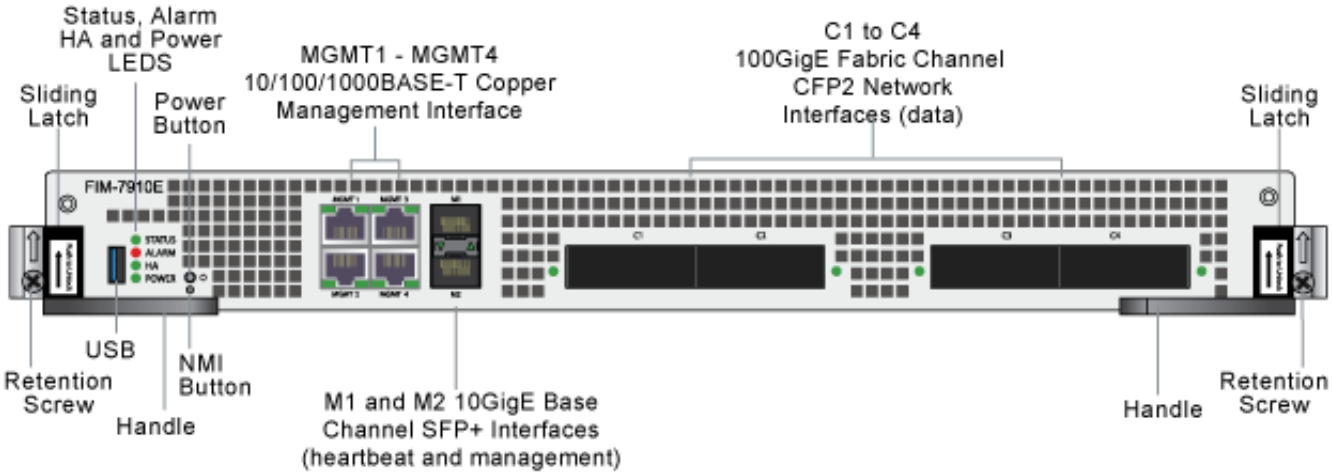

The FIM-7910E includes the following hardware features:

- Four front panel 100GigE CFP2 fabric channel interfaces (C1 to C4). These interfaces are connected to 100Gbps networks to distribute sessions to the FPM processor modules installed in chassis slots 3 and up. Using 100GBASE-SR10 multimode CFP2 transceivers, each CFP2 interface can also be split into ten 10GBASE-SR interfaces. These interfaces also support creating link aggregation groups (LAGs) that can include interfaces from both FIM-7910Es.
- Two front panel 10GigE SFP+ interfaces (M1 and M2) that connect to the base backplane channel. These interfaces are used for heartbeat, session sync, and management communication between FIM-7910Es in different chassis. These interfaces can also be configured to operate as Gigabit Ethernet interfaces using SFP transceivers, but should not normally be changed. If you use switches to connect these interfaces, the switch ports should be able to accept packets with a maximum frame size of at least 1526. The M1 and M2 interfaces need to be on different broadcast domains. If M1 and M2 are connected to the same switch, Q-in-Q must be enabled on the switch.
- Four 10/100/1000BASE-T out of band management Ethernet interfaces (MGMT1 to MGMT4).
- One 80Gbps fabric backplane channel for traffic distribution with each FPM installed in the same chassis as the FIM-7910E.
- One 1Gbps base backplane channel for base backplane with each FPM installed in the same chassis as the FIM-7910E.
- One 40Gbps fabric backplane channel for fabric backplane communication with the other FIM-7910E in the chassis.
- One 1Gbps base backplane channel for base backplane communication with the other FIM-7910E in the chassis.
- On-board DP2 processors and an integrated switch fabric to provide high-capacity session-aware load balancing.
- One front panel USB port.
- Power button.
- NMI switch (for troubleshooting as recommended by Fortinet Support).
- Mounting hardware.
- LED status indicators.

Splitting the FIM-7910E C1 to C4 interfaces

Each 100GE interface (C1 to C4) on the FIM-7910Es in slot 1 and slot 2 of a FortiGate-7000 system can be split into 10 x 10GBE interfaces. You split these interfaces after the FIM-7910Es are installed in your FortiGate-7000 system and the system us up and running. You can split the interfaces of the FIM-7910Es in slot 1 and slot 2 at the same time by entering a single CLI command. Splitting the interfaces requires a system reboot so Fortinet recommends that you split multiple interfaces at the same time according to your requirements to avoid traffic disruption.

For example, to split the C1 interface of the FIM-7910E in slot 1 (this interface is named 1-C1) and the C1 and C4 interfaces of the FIM-7910E in slot 2 (these interfaces are named 2-C1 and 2-C4) connect to the CLI of your FortiGate-7000 system using the management IP and enter the following command:

```
config system global
   set split-port 1-C1 2-C1 2-C4
end
```

After you enter the command, the FortiGate-7000 reboots and when it comes up:

- The 1-C1 interface will no longer be available. Instead the 1-C1/1, 1-C1/2, ..., and 1-C1/10 interfaces will be available.
- The 2-C1 interface will no longer be available. Instead the 2-C1/1, 2-C1/2, ..., and 2-C1/10 interfaces will be available.
- The 2-C4 interface will no longer be available. Instead the 2-C4/1, 2-C4/2, ..., and 2-C4/10 interfaces will be available.

You can now connect breakout cables to these interfaces and configure traffic between them just like any other FortiGate interface.

FIM-7910E hardware schematic

The FIM-7910E includes an integrated switch fabric (ISF) that connects the front panel interfaces to the DP2 session-aware load balancers and to the chassis backplanes. The ISF also allows the DP2 processors to distribute sessions among all NP6 processors on the FortiGate Processor Modules (FPM) in the FortiGate-7000chassis.

FIM-7910E hardware schematic

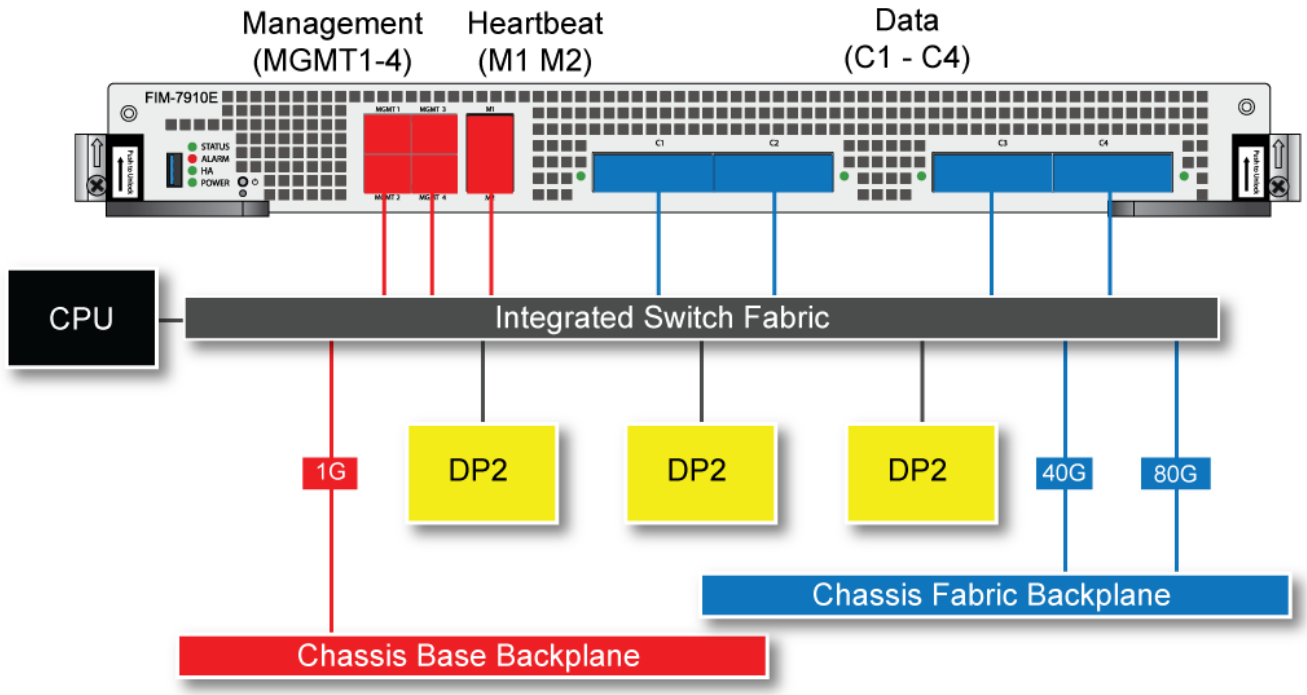

FIM-7920E interface module

The FIM-7920E interface module is a hot swappable module that provides data, management and session sync/heartbeat interfaces, base backplane switching and fabric backplane session-aware load balancing for a FortiGate-7000 series chassis. The FIM-7920E includes an integrated switch fabric and DP2 processors to load balance millions of data sessions over the chassis fabric backplane to FPM processor modules.

The FIM-7920E can be installed in any FortiGate-7000 series chassis in hub/switch slots 1 or 2. The FIM-7920E provides four Quad Small Form-factor Pluggable 28 (QSFP28) 100GigE interfaces for a FortiGate-7000 chassis. Using a 100GBASE-SR4 QSFP28 or 40GBASE-SR4 QSFP+ transceiver, each QSFP28 interface can also be split into four 10GBASE-SR interfaces.

You can also install FIM-7920Es in a second chassis and operate the chassis in HA mode with another set of processor modules to provide chassis failover protection.

FIM-7920E front panel

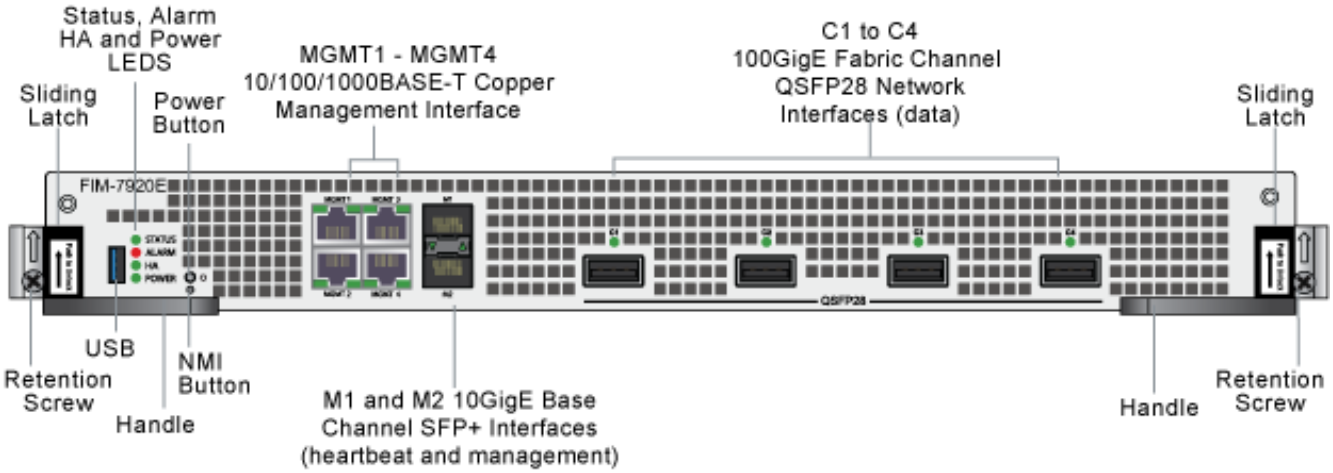

The FIM-7920E includes the following hardware features:

- Four front panel 100GigE QSFP28 fabric channel interfaces (C1 to C4). These interfaces are connected to 100Gbps networks to distribute sessions to the FPM processor modules installed in chassis slots 3 and up. Using a 100GBASE-SR4 QSFP28 or 40GBASE-SR4 QSFP+ transceiver, each QSFP28 interface can also be split into four 10GBASE-SR interfaces. These interfaces also support creating link aggregation groups (LAGs) that can include interfaces from both FIM-7920Es.
- Two front panel 10GigE SFP+ interfaces (M1 and M2) that connect to the base backplane channel. These interfaces are used for heartbeat, session sync, and management communication between FIM-7920Es in different chassis. These interfaces can also be configured to operate as Gigabit Ethernet interfaces using SFP transceivers, but should not normally be changed. If you use switches to connect these interfaces, the switch ports should be able to accept packets with a maximum frame size of at least 1526. The M1 and M2 interfaces need to be on different broadcast domains. If M1 and M2 are connected to the same switch, Q-in-Q must be enabled on the switch.
- Four 10/100/1000BASE-T out of band management Ethernet interfaces (MGMT1 to MGMT4).
- One 80Gbps fabric backplane channel for traffic distribution with each FPM installed in the same chassis as the FIM-7920E.
- One 1Gbps base backplane channel for base backplane with each FPM installed in the same chassis as the FIM-7920E.
- One 40Gbps fabric backplane channel for fabric backplane communication with the other FIM-7920E in the chassis.
- One 1Gbps base backplane channel for base backplane communication with the other FIM-7920E in the chassis.
- On-board DP2 processors and an integrated switch fabric to provide high-capacity session-aware load balancing.
- One front panel USB port.
- Power button.
- NMI switch (for troubleshooting as recommended by Fortinet Support).
- Mounting hardware.
- LED status indicators.

Changing the interface type and splitting the FIM-7920E C1 to C4 interfaces

By default, the FIM-7920E C1 to C4 interfaces are configured as 100GE QSFP28 interfaces. You can use the following command to convert them to 40GE QSFP+ interfaces. Once converted, you can use the other command below to split them into four 10GBASE-SR interfaces.

Changing the interface type

For example, to change the interface type of the C1 interface of the FIM-7920E in slot 1 to 40GE QSFP+ connect to the CLI of your FortiGate-7000 system using the management IP and enter the following command:

```
config system global
   set qsfp28-40g-port 1-C1
end
```

The FortiGate-7000 system reboots and when it starts up interface C1 of the FIM-7920E in slot 1 is operating as a 40GE QSFP+ interface .

To change the interface type of the C3 and C4 ports of the FIM-7920E in slot 2 to 40GE QSFP+ enter the following command:

```
config system global
   set qsfp28-40g-port 2-C3 2-C4
end
```

The FortiGate-7000 system reboots and when it starts up interfaces C3 and C4 of the FIM-7920E in slot 2 are operating as a 40GE QSFP+ interfaces.

Splitting the C1 to C4 interfaces

Each 40GE interface (C1 to C4) on the FIM-7920Es in slot 1 and slot 2 of a FortiGate-7000 system can be split into 4 x 10GBE interfaces. You split these interfaces after the FIM-7920Es are installed in your FortiGate-7000 system and the system us up and running. You can split the interfaces of the FIM-7920Es in slot 1 and slot 2 at the same time by entering a single CLI command. Splitting the interfaces requires a system reboot so Fortinet recommends that you split multiple interfaces at the same time according to your requirements to avoid traffic disruption.

For example, to split the C1 interface of the FIM-7920E in slot 1 (this interface is named 1-C1) and the C1 and C4 interfaces of the FIM-7920E in slot 2 (these interfaces are named 2-C1 and 2-C4) connect to the CLI of your FortiGate-7000 system using the management IP and enter the following command:

```
config system global
   set split-port 1-C1 2-C1 2-C4
end
```

After you enter the command, the FortiGate-7000 reboots and when it comes up:

- The 1-C1 interface will no longer be available. Instead the 1-C1/1, 1-C1/2, 1-C1/3, and 1-C1/4 interfaces will be available.
- The 2-C1 interface will no longer be available. Instead the 2-C1/1, 2-C1/2, 2-C1/3, and 2-C1/4 interfaces will be available.
- The 2-C4 interface will no longer be available. Instead the 2-C4/1, 2-C4/2, 2-C4/3, and 2-C4/4 interfaces will be available.

You can now connect breakout cables to these interfaces and configure traffic between them just like any other FortiGate interface.

FIM-7920E hardware schematic

The FIM-7920E includes an integrated switch fabric (ISF) that connects the front panel interfaces to the DP2 session-aware load balancers and to the chassis backplanes. The ISF also allows the DP2 processors to distribute sessions among all NP6 processors on the FortiGate Processor Modules (FPM) in the FortiGate-7000 chassis.

FIM-7920E hardware schematic

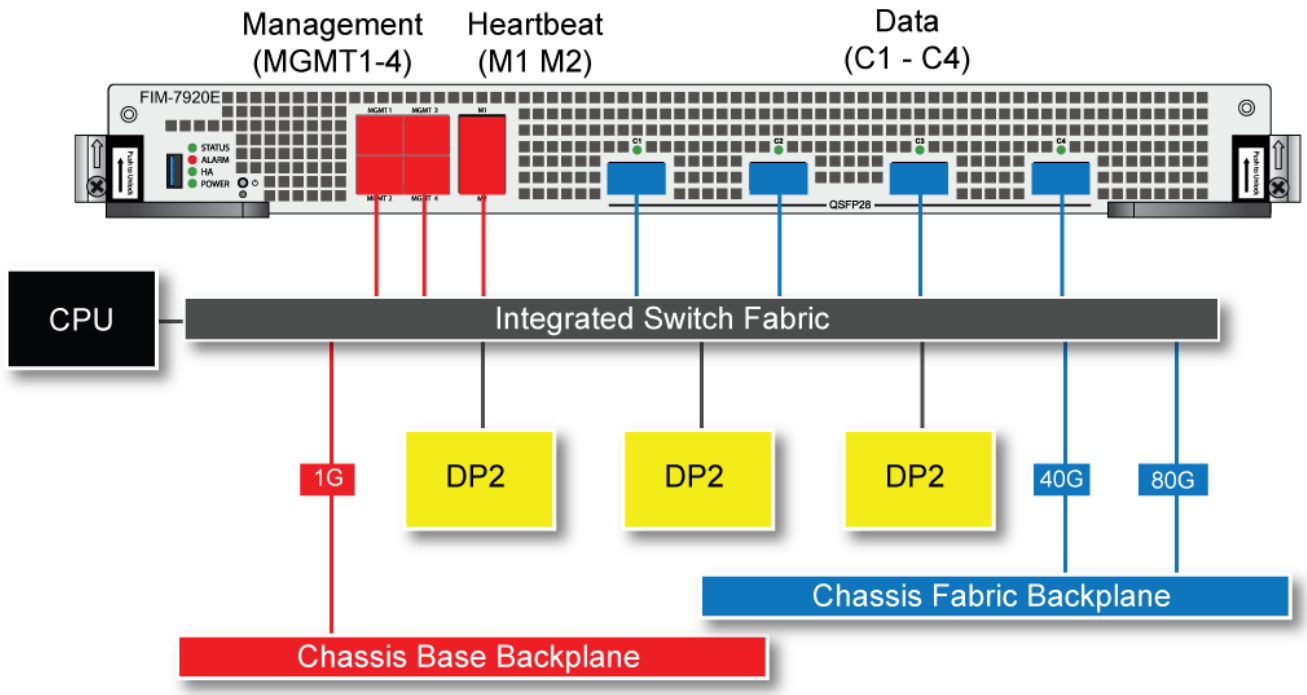

FPM-7620E processing module

The FPM-7620E processing module is a high-performance worker module that processes sessions load balanced to it by FortiGate-7000 series interface (FIM) modules over the chassis fabric backplane. The FPM-7620E can be installed in any FortiGate-7000 series chassis in slots 3 and up.

The FPM-7620E includes two 80Gbps connections to the chassis fabric backplane and two 1Gbps connections to the base backplane. The FPM-7620E processes sessions using a dual CPU configuration, accelerates network traffic processing with 4 NP6 processors and accelerates content processing with 8 CP9 processors. The NP6 network processors are connected by the FIM switch fabric so all supported traffic types can be fast path accelerated by the NP6 processors.

The FPM-7620E includes the following hardware features:

- Two 80Gbps fabric backplane channels for load balanced sessions from the FIMs installed in the chassis.
- Two 1Gbps base backplane channels for management, heartbeat and session sync communication.
- Dual CPUs for high performance operation.
- Four NP6 processors to offload network processing from the CPUs.
- Eight CP9 processors to offload content processing and SSL and IPsec encryption from the CPUs.

FPM-7620E front panel

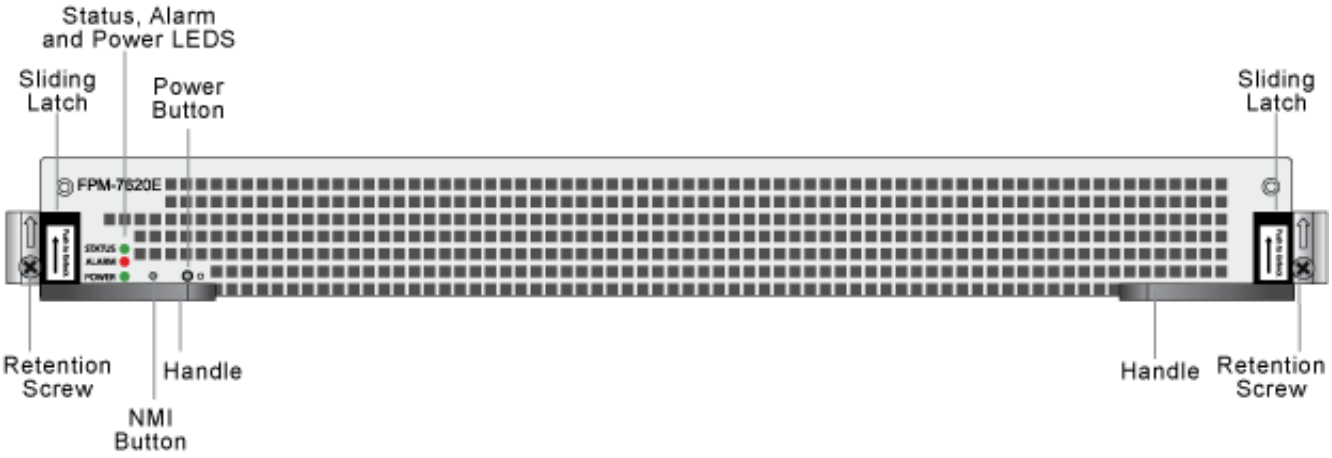

- Power button.
- NMI switch (for troubleshooting as recommended by Fortinet Support).
- Mounting hardware.
- LED status indicators.

NP6 network processors - offloading load balancing and network traffic

In a FortiGate-7000 chassis, FPM-7620E NP6 network processors combined with the FortiGate Interface Module (FIM) Integrated Switch Fabric (ISF) provide hardware acceleration by offloading sessions from the FPM-7620E CPUs. The result is enhanced network performance provided by the NP6 processors plus the removal of network processing load from the FPM-7620 CPUs. The NP6 processors can also handle some CPU-intensive tasks, like IPsec VPN encryption/decryption. Because of the ISF in each FIM, all sessions are fast-pathed and accelerated.

FPM-7620E hardware architecture

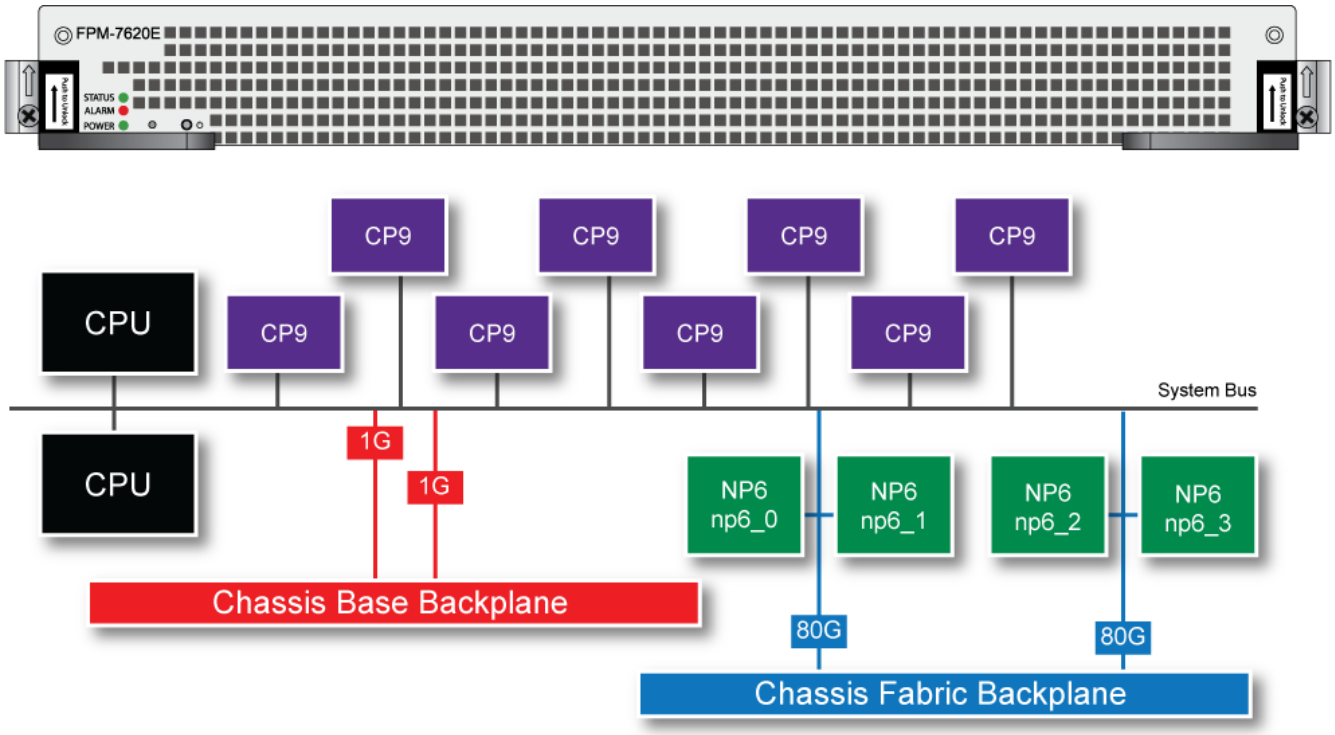

Accelerated IPS, SSL VPN, and IPsec VPN (CP9 content processors)

The FPM-7620E includes eight CP9 processors that provide the following performance enhancements:

- Flow-based inspection (IPS, application control etc.) pattern matching acceleration with over 10Gbps throughput
- IPS pre-scan
- IPS signature correlation
- Full match processors
- High performance VPN bulk data engine
- IPsec and SSL/TLS protocol processor
- DES/3DES/AES128/192/256 in accordance with FIPS46-3/FIPS81/FIPS197
- MD5/SHA-1/SHA256/384/512-96/128/192/256 with RFC1321 and FIPS180
- HMAC in accordance with RFC2104/2403/2404 and FIPS198
- ESN mode
- GCM support for NSA "Suite B" (RFC6379/RFC6460) including GCM-128/256; GMAC-128/256
- Key Exchange Processor that supports high performance IKE and RSA computation
- Public key exponentiation engine with hardware CRT support
- Primary checking for RSA key generation
- Handshake accelerator with automatic key material generation
- True Random Number generator
- Elliptic Curve support for NSA "Suite B"

- Sub public key engine (PKCE) to support up to 4096 bit operation directly (4k for DH and 8k for RSA with CRT)
- DLP fingerprint support
- TTTD (Two-Thresholds-Two-Divisors) content chunking
- Two thresholds and two divisors are configurable

Getting started with FortiGate-7000

Once you have installed your FortiGate-7000 chassis in a rack and installed FIM interface modules and FPM processing modules in it you can power on the chassis and all modules in the chassis will power up.

Whenever a chassis is first powered on, it takes about 5 minutes for all modules to start up and become completely initialized and synchronized. During this time the chassis will not allow traffic to pass through and you may not be able to log into the GUI, or if you manage to log in the session could time out as the FortiGate-7000 continues negotiating.

Review the chassis and module front panel LEDs to verify that everything is operating normally. Wait until the chassis has complete started up and synchronized before making configuration changes.

You can use the `diagnose sys ha status` command to confirm that the FortiGate-7000 is completely initialized. If the output from entering this command hasn't changed after checking for a few minutes you can assume that the system has initialized. You don't normally have to confirm that the system has initialized, but this diagnose command is available to verify that the system is initialized.

Once the system has initialized you have a few options for connecting to the FortiGate-7000 GUI or CLI:

- Log into the GUI by connecting the MGMT1 interface of the FIM in slot 1 to your network. Then browse to https://192.168.1.99. The default administrator account is **admin** with no password. For more information about using the MGMT interfaces for management connections to the FortiGate-7000, see Setting up management connections on page 36.
- Log into the CLI by connecting the MGMT1 interface of the FIM in slot 1 to your network. Then use an SSH client to connect to 192.168.1.99 and use the same admin account to log in.
- Log into the CLI with a serial connection to the Console 1 serial port on the FortiGate-7000 management module. For more information, including the serial port settings, see Connecting to module CLIs using the management module on page 40.

The following example Unit Operation dashboard widget shows a FortiGate-7040E with the following modules:

- An FIM-7904E in slot 1. Each of the FIM-7904E B1 to B8 front panel interfaces has been split into four interfaces
- An FIM-7910E in slot 2. Each of the FIM0-7910E C1 to C4 front panel interfaces has been split into ten interfaces.
- FPM-7620Es in slots 3 and 4.

Example FortiGate-7040 unit operation widget view

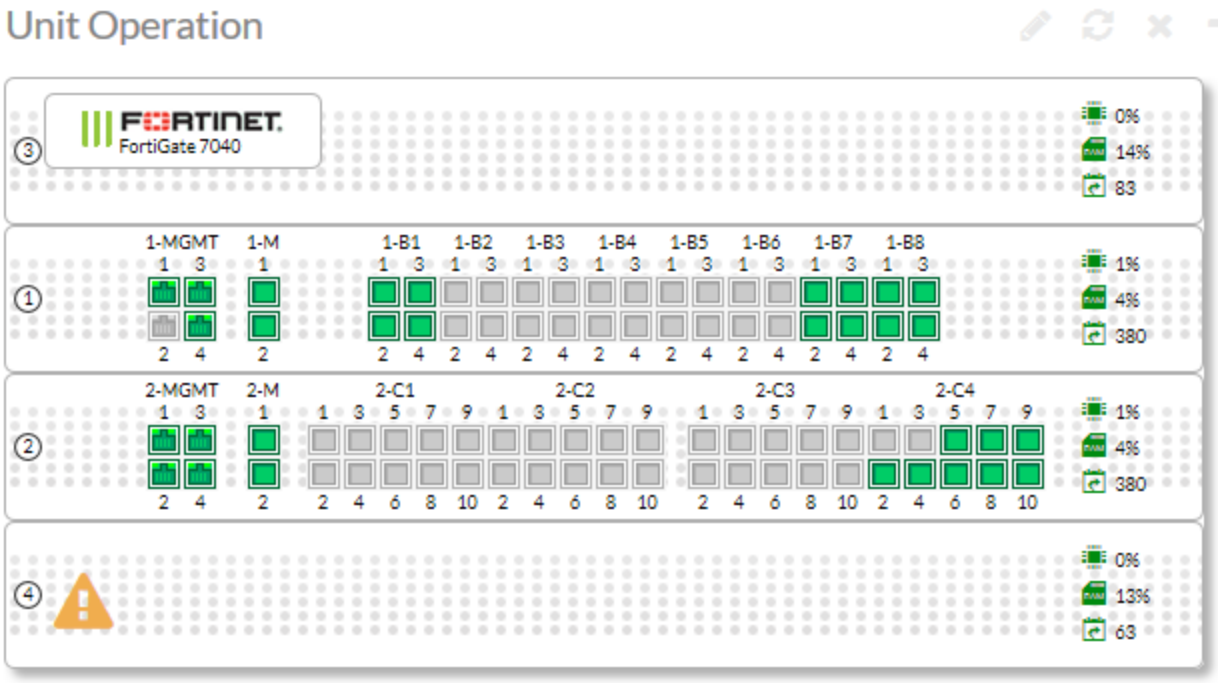

Setting up management connections

When your FortiGate-7000 first starts up, the MGMT1 to MGMT4 ports of both of the FIMs are part of a static 802.3 aggregate interface with a default IP address of 192.168.1.99. On the GUI or CLI the 802.3 aggregate interface is named **mgmt**.

Example mgmt management interface configuration

Interface Name	mgmt
Link Status	Up
Type	802.3ad Aggregate
Virtual Domain	dmgmt-vdom
Physical Interface Members	1-mgmt1, 1-mgmt2, 1-mgmt3, 1-mgmt4, 2-mgmt1, 2-mgmt2, 2-mgmt3, 2-mgmt4

Setting up a single management connection

You can configure and manage your FortiGate-7000 by connecting an Ethernet cable to any of the MGMT1 - 4 interfaces of the FIM in slot 1 or slot 2 and logging into the GUI using HTTPS or the CLI using SSH. The default IP address is 192.168.1.99 and you can log in with the **admin** administrator account with no password.

For security reasons you should add a password to the admin account before connecting the chassis to your network.

Setting up redundant management connections

You can set up redundant management connections to your FortiGate-7000 by adding a static 802.3 aggregate interface to a switch and setting up multiple connections between the switch and the FIM MGMT ports. Then connect the switch to your network.

LACP is not supported for the mgmt aggregate interface.

You do not have the change the configuration of the FortiGate-7000 to set up redundant management connections. The following example shows connections between the MGMT 1 interfaces of each FIM to a switch. The switch is configured with a 802.3 static aggregate interface that includes two ports, one for each MGMT 1 interface. The switch also connects the MGMT 1 interfaces to a management network.

Example FortiGate-7000 redundant management connections

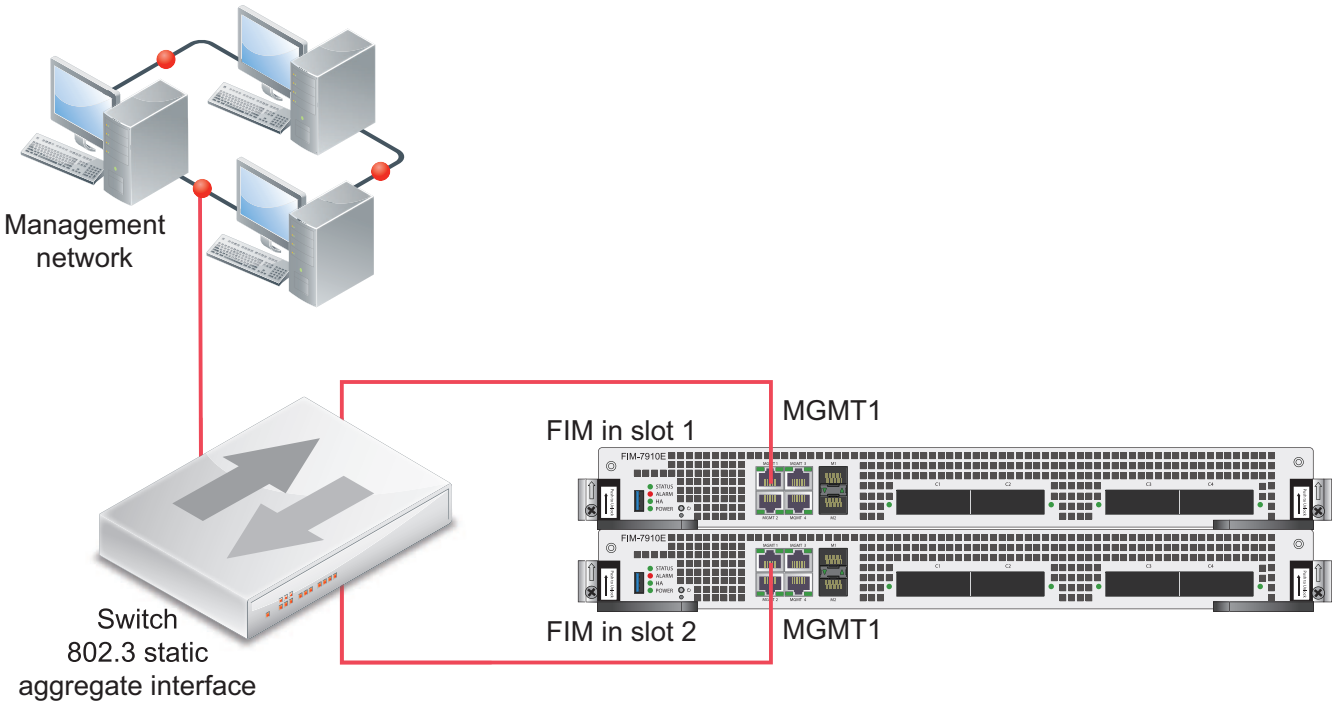

The following example shows redundant connections between both FIMs and the switch. In this case you need to add more switch ports to the static aggregate interface on the switch. You do not have to change the configuration of the FortiGate-7000 to set up this redundant management connection configuration.

Example FortiGate-7000 redundant management connections with redundant connections to each FIM

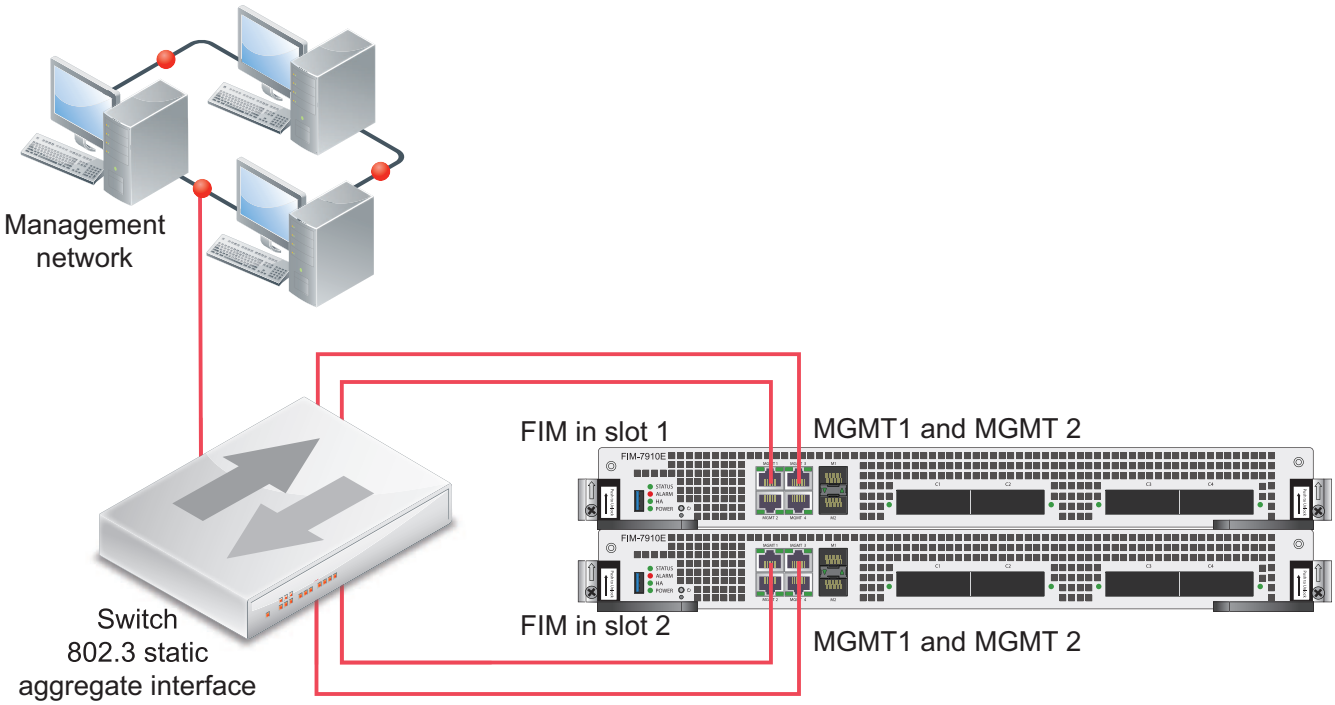

In either of these configurations, for additional redundancy you can use two switches. If you use two redundant switches, the static aggregate interface should span across both switches.

Managing individual FIMs and FPMs

When you log into the GUI or CLI using the mgmt interface IP address you are actually connected to the primary (or master) FIM in slot 1 (the address of slot 1 is FIM01). To verify which module you have logged into, the GUI header banner or CLI prompt shows the hostname of the module you are logged into plus the slot address in the format <hostname> (<slot address>).

In some cases you may want to connect to individual modules. For example, you may want to view the traffic being processed by a specific FPM. You can connect to the GUI or CLI of individual modules in the chassis using the system management IP address with a special port number.

For example, if the system management IP address is 192.168.1.99 you can connect to the GUI of the FIM in slot 1 using the system management IP address (for example, by browsing to https://192.168.1.99). You can also use the system management IP address followed by the special port number, for example 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 chassis slot number (in this example, 01). The following table lists the special ports to use to connect to each chassis slot using common admin protocols:

FortiGate-7000 special administration port numbers

Slot Number	Slot Address	HTTP (80)	HTTPS (443)	Telnet (23)	SSH (22)	SNMP (161)
5	FPM05	8005	44305	2305	2205	16105
3	FPM03	8003	44303	2303	2203	16103
1	FIM01	8001	44301	2301	2201	16101
2	FIM02	8002	44302	2302	2202	16102
4	FPM04	8004	44304	2304	2204	16104
6	FPM06	8006	44306	2306	2206	16106

For example:

- To connect to the GUI of the FIM in slot 3 using HTTPS you would browse to https://192.168.1.99:44303.
- To send an SNMP query to the FPM in slot 6 use the port number 16106.

The FortiGate-7000 configuration is the same no matter which modem you log into. Logging into different modules allows you to use FortiView or Monitor GUI pages to view the activity on that module. Even though you can log into different modules, you should only make configuration changes from the primary FIM; which is the FIM in slot 1.

Managing individual modules from the CLI

From the CLI you can use the following command to switch between chassis slots and perform different operations on the modules in each slot:

```
execute load-balance slot {manage | power-off | power-on | reboot} <slot-number>
```

Use `manage` to connect to the CLI of a different module, use `power-off`, `power-on`, and `reboot` to turn off or turn on the power or reboot the module in `<slot-number>`.

Connecting to module CLIs using the management module

All FortiGate-7000 chassis includes a front panel management module (also called a shelf manager) on the chassis front panel. See the system guide for your chassis for details about the management module.

ForiGate-7040E management module front panel

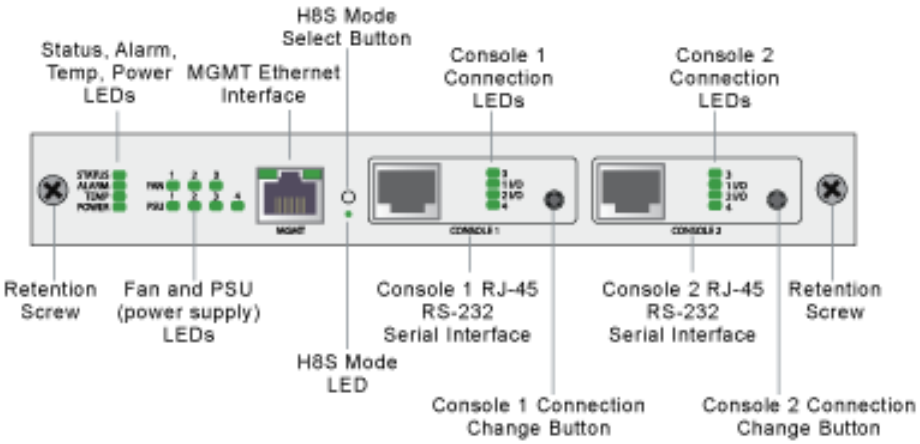

The management module includes two console ports named Console 1 and Console 2 that can be used to connect to the CLI of the FIM and FPMs in the chassis. As described in the system guide, the console ports are also used to connect to SMC CLIs of the management module and the FIM and FPMs

By default when the chassis first starts up Console 1 is connected to the FortiOS CLI of the FIM in slot 1 and Console 2 is disconnected. The default settings for connecting to each console port are:

Baud Rate (bps) 9600, Data bits 8, Parity None, Stop bits 1, and Flow Control None.

You can use the console connection change buttons to select the CLI that each console port is connected to. Press the button to cycle through the FIM and FPM FortiOS CLIs and disconnect this console. The console's LEDs indicate what it is connected to. If no LED is lit the console is either connected to the management module SMC SDI console or disconnected. Both console ports cannot be connected to the same CLI at the same time. If a console button press would cause a conflict that module is skipped. If one of the console ports is disconnected then the other console port can connect to any CLI.

If you connect a PC to one of the management module console ports with a serial cable and open a terminal session you begin by pressing Ctrl-T to enable console switching mode. Press Ctrl-T multiple times to cycle through the FIM and FPM FortiOS CLIs (the new destination is displayed in the terminal window). If you press Ctrl-T after connecting to the FPM in the highest slot number, the console is disconnected. Press Ctrl-T again to start over again at slot 1.

Once the console port is connected to the CLI that you want to use, press Enter to enable the CLI and login. The default administrator account for accessing the FortiOS CLIs is admin with no password.

When your session is complete you can press Ctrl-T until the prompt shows you have disconnected from the console.

Connecting to the FortiOS CLI of the FIM in slot 1

Use the following steps to connect to the FortiOS CLI of the FIM in slot 1:

1. Connect the console cable supplied with your chassis to Console 1 and to your PC or other device RS-232 console port.
2. 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.
3. Press Ctrl-T to enter console switch mode.
4. Repeat pressing Ctrl-T until you have connected to slot 1. Example prompt:

```
<Switching to Console: FIM01 (9600)>
```

5. Login with an administrator name and password.

   The default is admin with no password. For security reasons, it is strongly recommended that you change the password.
6. When your session is complete, enter the `exit` command to log out.

Default VDOM configuration

By default when the FortiGate-7000 first starts up it is operating in multiple VDOM mode. The system has a management VDOM (named dmgmt-vdom) and the root VDOM. All management interfaces are in dmgmt-vdom and all other interfaces are in the root VDOM. You can add more VDOMs and add interfaces to them or just use the root VDOM.

Default management VDOM

By default the FortiGate-7000 configuration includes a management VDOM named dmgmt-vdom. For the FortiGate-7000 system to operate normally you should not change the configuration of this VDOM and this VDOM should always be the management VDOM. You should also not add or remove interfaces from this VDOM.

You have full control over the configurations of other FortiGate-7000 VDOMs.

Firmware upgrades

All of the modules in your FortiGate-7000 run the same firmware image. You upgrade the firmware from the primary FIM GUI or CLI just as you would any FortiGate product. During the upgrade process the firmware of all of the modules in the chassis upgrades in one step. Firmware upgrades should be done during a quiet time because traffic will briefly be interrupted during the upgrade process.

Before beginning a firmware upgrade, Fortinet recommends the following:

- 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.
- Backup your FortiGate-7000 HA configuration.
- Review the services provided by your FortiGate-7000 before the upgrade and then again after the upgrade to make sure everything continues to operate normally. For example, you might want to verify that you can successfully access a key server used by your organization before the upgrade and make sure after the upgrade that you can still reach the server and that performance is comparable. You could also take a snapshot of key performance indicators (number of sessions, CPU usage, and memory usage) before the upgrade and verify that you see comparable performance after the upgrade .

For information about upgrading FortiGate-7000 HA configuration, see Upgrading FortiGate-7000 HA cluster firmware on page 72.

Restarting the FortiGate-7000

To restart all of the modules in a FortiGate-7000 chassis, connect to the primary FIM CLI and enter the command `execute reboot`. When you enter this command all of the modules in the chassis reboot.

You can restart individual modules by logging into that module's CLI and entering the `execute reboot` command.

Configuration synchronization

The FortiGate-7000 synchronizes the configuration to all modules in the chassis. To support this feature, the FIM in slot 1 becomes the config-sync master and this FIM makes sure the configurations of all modules are synchronized. Every time you make a configuration change you must be logged into the chassis using the management address, which logs you into the config-sync master. All configuration changes made to the config-sync master are synchronized to all of the modules in the chassis.

If the FIM in slot 1 fails or reboots, the FIM in slot 2 becomes the config-sync master.

Packet sniffing for FIM and FPM packets

You can use the `diagnose sniffer packet` command to view or sniff packets as they are processed by FIM or FPMs. To use this command you have to be logged into a VDOM. You can run this command from any FIM or FPM CLI.

If you run the command from the primary FIM you can use the `<slot>` option to view packets for the module in that slot. If you run the command from an individual FIM or FPM the `<slot>` option is not available and you see only the packets processed by that module.

From an FIM the command syntax is:

```
diagnose sniffer packet <interface> <protocol-filter> <verbose> <count> <timestamp> <slot>
```

Where:

`<interface>` is the name of one or more interfaces on which to sniff for packets. Use `any` to sniff packets for all interfaces. To view management traffic use the `elbc-base-ctrl` interface name.

`<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.

When you press Enter you are prompted to run the packet capture on the module that you have logged into or you can input a slot number (for example, `slot2`) to capture packets on the module in that slot.

Load balancing and flow rules

This chapter provides an overview of how FortiGate-7000 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-7000 SLBC works as follows.

1. SLBC attempts to match all incoming sessions with a configured flow rule (see Load balancing and flow rules on page 45). If a session matches a flow rule, the session is directed according to the action setting of the flow rule. Usually flow rules send traffic that can't be load balanced to a specific FPM.
2. TCP, UDP, SCTP, ICMP (IPv4 only) and ESP (IPv4 only) sessions that do not match a flow rule are directed to the DP2 processors.
   The DP2 processors distribute sessions to the FPMs according to the load balancing method set by the `dp-load-distribution-method` option of the `config load-balance setting` command.
3. All other sessions are sent to the primary FPM.

Setting the load balancing method

Sessions are load balanced or distributed 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
```

Where:

`to-master` directs all session to the primary FPM. This method is for troubleshooting only and should not be used for normal operation. Directing all sessions to the primary FPM will have a negative impact on performance.

`src-ip` sessions are distributed across all FPMs according to their source IP address.

`dst-ip` sessions are statically distributed across all FPMs according to their destination IP address.

`src-dst-ip` sessions are distributed across all FPMs according to their source and destination IP addresses.

`src-ip-sport` sessions are distributed across all FPMs according to their source IP address and source port.

`dst-ip-dport` sessions are distributed across all FPMs according to their destination IP address and destination port.

`src-dst-ipsport-dport` sessions are distributed across all FPMs 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. All session information is taken 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.

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 (or master) FPM by configuring flow rules for that traffic. You can also configure flow rules to send traffic that cannot be load balanced to specific FPMs.

Create flow rules using the `config loadbalance flow-rule` command. The default configuration uses this command to send IKE, GRE, session helper, Kerberos, BGP, RIP, IPv4 and IPv6 DHCP, PPTP, BFD, IPv4 multicast and IPv6 multicast to the primary FPM. You can view the default configuration of the `config loadbalance flow-rule` command to see how this is all configured. For example, the following configuration sends all IKE sessions to the primary FPM:

```
config load-balance flow-rule
   edit 1
      set status enable
      set vlan 0
      set ether-type ip
      set protocol udp
      set src-l4port 500-500
      set dst-l4port 500-500
      set action forward
      set forward-slot master
      set priority 5
      set comment "ike"
   next
   edit 2
      set status disable
      set vlan 0
      set ether-type ip
      set protocol udp
      set src-l4port 4500-4500
      set dst-l4port 0-0
      set action forward
      set forward-slot master
      set priority 5
      set comment "ike-natt src"
   next
   edit 3
      set status disable
      set vlan 0
      set ether-type ip
      set protocol udp
      set src-l4port 0-0
      set dst-l4port 4500-4500
      set action forward
      set forward-slot master
      set priority 5
      set comment "ike-natt dst"
   next
```

See Default configuration for traffic that cannot be load balanced on page 49 for a listing of all of the default configuration flow rules.

Determining the primary FPM

You can use the `diagnose load-balance status` command to determine which FPM is operating as the primary FPM. The primary FPM is also responsible for dynamic routing.

Example `diagnose load-balance status` output for a FortiGate-7060E showing that the FPM in slot 3 is the primary (master) FPM. The command output also shows the status of all of the FPMs in the FortiGate-7060E. The output also shows that the FPM in slot 4 is either missing or down.

```
diagnose load-balance status
============================================================================
Slot: 2  Module SN: FIM04E3E16000222
  FIM02: FIM04E3E16000222
  Master FPM Blade: slot-3

     Slot  3: FPM20E3E17900133
       Status:Working   Function:Active
       Link:      Base: Up          Fabric: Up
       Heartbeat: Management: Good   Data: Good
       Status Message:"Running"
     Slot  4:
       Status:Dead      Function:Active
       Link:      Base: Up          Fabric: Down
       Heartbeat: Management: Failed Data: Failed
       Status Message:"Waiting for management heartbeat."
.

.

.
```

IPsec VPN load balancing

You can use the following command to enable or disable IPsec VPN load balancing:

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

By default IPsec VPN load balancing is enabled and the flow rules listed below are disabled. So the IPsec VPN sessions is directed to the DP2 processors and load balanced to the FPMs.

However, IPsec VPN load balancing enabled, IPsec VPN sessions traveling between two IPsec tunnels will be dropped because the two IPsec tunnels may be terminated on different FPMs. So if you have traffic entering the FortiGate-7000 from one IPsec VPN tunnel and leaving the FortiGate-7000 out another IPsec VPN tunnel you need to disable IPsec load balancing:

```
config load-balance setting
   config ipsec-load-balance disable
end
```

The following flow rules are enabled if IPsec VPN load balancing is disabled:

```
config load-balance flow-rule
    edit 22
        set ether-type ipv4
        set protocol udp
        set src-l4port 500-500
        set dst-l4port 500-500
        set comment "ipv4 ike"
    next
    edit 23
        set ether-type ipv4
        set protocol udp
        set src-l4port 4500-4500
        set comment "ipv4 ike-natt src"
    next
    edit 24
        set ether-type ipv4
        set protocol udp
        set dst-l4port 4500-4500
        set comment "ipv4 ike-natt dst"
    next
    edit 25
        set ether-type ipv4
        set protocol esp
        set comment "ipv4 esp"
    next
end
```

These flow rules should generally handle all IPsec VPN traffic. You can also adjust them or add your own flow rules if you have an IPsec VPN setup that is not compatible with the default flow rules.

GTP load balancing

You can use the following command to enable or disable GTP load balancing.

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

By default, this option is disabled and GTP load balancing is disabled. The following flow rule is enabled and GTP traffic is directed to the primary (master) FPM.

```
config load-balance flow-rule
    edit 17
        set ether-type ipv4
        set protocol udp
        set dst-l4port 2123-2123
        set comment "gtp-c to master blade"
    next
end
```

When the `gtp-load-balance` option is enabled, the GTP load balancing is flow rule is disabled and GTP sessions are directed to the DP2 processors and load balanced to the FPMs.

Default configuration for traffic that cannot be load balanced

The following flow rules are recommended to handle common forms of traffic that cannot be load balanced. These flow rules send GPRS (port 2123), SSL VPN, IPv4 and IPv6 IPsec VPN, ICMP and ICMPv6 traffic to the primary (or master) FPM.

The CLI syntax below just shows the configuration changes. All other options are set to their defaults. For example, the flow rule option that controls the FPM slot that sessions are sent to is `forward-slot` and in all cases below `forward-slot` is set to its default setting of `master`. This setting sends matching sessions to the primary (or master) FPM.

```
config load-balance flow-rule
   edit 20
      set status enable
      set ether-type ipv4
      set protocol udp
      set dst-l4port 2123-2123
   next
   edit 21
      set status enable
      set ether-type ip
      set protocol tcp
      set dst-l4port 10443-10443
      set comment "ssl vpn to the primary FPM"
   next
   edit 22
      set status enable
      set ether-type ipv4
      set protocol udp
      set src-l4port 500-500
      set dst-l4port 500-500
      set comment "ipv4 ike"
   next
   edit 23
      set status enable
      set ether-type ipv4
      set protocol udp
      set src-l4port 4500-4500
      set comment "ipv4 ike-natt src"
   next
   edit 24
      set status enable
      set ether-type ipv4
      set protocol udp
      set dst-l4port 4500-4500
      set comment "ipv4 ike-natt dst"
   next
   edit 25
      set status enable
      set ether-type ipv4
      set protocol esp
      set comment "ipv4 esp"
   next
   edit 26
      set status enable
      set ether-type ipv6
```

```
      set protocol udp
      set src-l4port 500-500
      set dst-l4port 500-500
      set comment "ipv6 ike"
   next
   edit 27
      set status enable
      set ether-type ipv6
      set protocol udp
      set src-l4port 4500-4500
      set comment "ipv6 ike-natt src"
   next
   edit 28
      set status enable
      set ether-type ipv6
      set protocol udp
      set dst-l4port 4500-4500
      set comment "ipv6 ike-natt dst"
   next
   edit 29
      set status enable
      set ether-type ipv6
      set protocol esp
      set comment "ipv6 esp"
   next
   edit 30
      set ether-type ipv4
      set protocol icmp
      set comment "icmp"
   next
   edit 31
      set status enable
      set ether-type ipv6
      set protocol icmpv6
      set comment "icmpv6"
   next
   edit 32
      set ether-type ipv6
      set protocol 41
   end
```

Operating a FortiGate-7000

This chapter describes some FortiGate-7000 general operating procedure.

Failover in a standalone FortiGate-7000

A FortiGate-7000 will continue to operate even if one of the FIM or FPMs fails or is removed. If an FPM fails, sessions being processed by that module fail. All sessions are then load balanced to the remaining FPMs. Sessions that were being processed by the failed module are restarted and load balanced to the remaining FPMs.

If an FIM fails, the other FIM will continue to operate and will become the config-sync master. However, traffic received by the failed FIM will be lost.

You can use LACP or redundant interfaces to connect interfaces of both FIMs to the same network. In this way, if one of the FIMs fails the traffic will continue to be received by the other FIM.

Replacing a failed FPM or FIM

This section describes how to remove a failed FPM or FIM and replace it with a new one. The procedure is slightly different depending on if you are operating in HA mode with two chassis or just operating a standalone chassis.

Replacing a failed module in a standalone FortiGate-7000 chassis

1. Power down the failed module by pressing the front panel power button.
2. Remove the module from the chassis.
3. Insert the replacement module. It should power up when inserted into the chassis if the chassis has power.
4. The module's configuration is synchronized and its firmware is upgraded to match the firmware version on the primary module. The new module reboots.
5. Confirm that the new module is running the correct firmware version either from the GUI or by using the `config system status` command.
   Manually update the module to the correct version if required. You can do this by logging into the module and performing a firmware upgrade.
6. Verify that the configuration has been synchronized.
   The following command output shows the sync status of the FIMs in a FortiGate-7000 chassis. The field `in_sync=1` indicates that the configurations of the modules are synchronized.

```
diagnose sys confsync
status | grep in_sy
FIM04E3E16000080, Slave, uptime=177426.45, priority=2,
slot_id=1:2, idx=0, flag=0x0, in_sync=1
FIM10E3E16000063, Master, uptime=177415.38, priority=1,
slot_id=1:1, idx=1, flag=0x0, in_sync=1
```

If `in_sync` is not equal to 1 or if a module is missing in the command output you can try restarting the modules in the chassis by entering execute reboot from any module CLI. If this does not solve the problem, contact Fortinet support.

Replacing a failed module in a FortiGate-7000 chassis in an HA cluster

1. Power down the failed module by pressing the front panel power button.
2. Remove the module from the chassis.
3. Insert the replacement module. It should power up when inserted into the chassis if the chassis has power.
4. The module's configuration is synchronized and its firmware is upgraded to match the configuration and firmware version on the primary module. The new module reboots.
5. Confirm that the module is running the correct firmware version.
   Manually update the module to the correct version if required. You can do this by logging into the module and performing a firmware upgrade.
6. Configure the new module for HA operation. For example:
   ```
   config system ha
      set mode a-p
      set chassis-id 1
      set hbdev m1 m2
      set hbdev-vlan-id 999
      set hbdev-second-vlan-id 990
   end
   ```
7. Optionally configure the hostname:
   ```
   config system global
      set hostname <name>
   end
   ```
   The HA configuration and the hostname must be set manually because HA settings and the hostname is not synchronized.
8. Verify that the configuration has been synchronized.
   The following command output shows the sync status of the FIMs in a FortiGate-7000 chassis. The field `in_sync=1` indicates that the configurations of the modules are synchronized.
   ```
   diagnose sys confsync
   status | grep in_sy
   FIM04E3E16000080, Slave, uptime=177426.45, priority=2,
   slot_id=1:2, idx=0, flag=0x0, in_sync=1
   FIM10E3E16000063, Master, uptime=177415.38, priority=1,
   slot_id=1:1, idx=1, flag=0x0, in_sync=1
   ```
   If `in_sync` is not equal to 1 or if a module is missing in the command output you can try restarting the modules in the chassis by entering execute reboot from any module CLI. If this does not solve the problem, contact Fortinet support.

Installing firmware on an FIM or FPM from the BIOS using a TFTP server

Use the procedures in this section to install firmware on a FIM or FPM from a TFTP server after interrupting the boot up sequence from the BIOS.

You might want to use this procedure if you need to reset the configuration of a module to factory defaults by installing firmware from a reboot. You can also use this procedure if you have formatted one or more FIM or FPMs from the BIOS by interrupting the boot process.

This procedure involves creating a connection between a TFTP server and one of the MGMT interfaces of one of the FIMs, using a chassis console port to connect to the CLI of the module that you are upgrading the firmware for, rebooting this module, interrupting the boot from the console session, and installing the firmware.

This section includes two procedures, one for upgrading FIMs and one for upgrading FPMs. The two procedures are very similar but a few details, most notably the local VLAN ID setting are different. If you need to update both FIM and FPMs, you should update the FIMs first as the FPMs can only communicate with the TFTP server through FIM interfaces.

Uploading firmware from a TFTP server to an FIM

Use the following steps to upload firmware from a TFTP server to an FIM. This procedure requires Ethernet connectivity between the TFTP server and one of the FIM's MGMT interfaces.

During this procedure, the FIM will not be able to process traffic so, if possible, perform this procedure when the network is not processing any traffic.

If you are operating an HA configuration, you should remove the FortiGate-7000 from the HA configuration before performing this procedure.

1. Set up a TFTP server and copy the firmware file to be installed into the TFTP server default folder.
2. Set up your network to allow traffic between the TFTP server and one of the MGMT interfaces of the FIM to be updated.

   If the MGMT interface you are using is one of the MGMT interfaces connected as a LAG to a switch you must shutdown or disconnect all of the other connections in the LAG from the switch. This includes the MGMT interfaces in the other FIM.
3. Connect the console cable supplied with your chassis to the Console 1 port on your chassis front panel and to your management computer's RS-232 console port.
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. Press Ctrl-T to enter console switch mode.
6. Repeat pressing Ctrl-T until you have connected to the module to be updated. Example prompt:
   `<Switching to Console: FIM02 (9600)>`
7. Optionally log into the FIM's CLI.
8. Reboot the FIM to be updated.

   You can do this using the `execute reboot` command from the CLI or by pressing the power switch on the module front panel.
9. When the FIM starts up, follow the boot process in the terminal session and press any key when prompted to interrupt the boot process.
10. Press `C` to set up the TFTP configuration.
11. Use the BIOS menu to set the following. Only change settings if required.

    `[P]: Set image download port`: MGMT1 (change if required)

    `[D]: Set DHCP mode`: Disabled

    `[I]: Set local IP address`: A temporary IP address to be used to connect to the TFTP server. This address must not be the same as the chassis 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`: Use -1 to clear the Local VLAN ID.

`[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 file to be installed.

12. Press Q to quit this menu.
13. Press R to review the configuration.

    If you need to make any corrections, press C and make the changes as required. When the configuration is correct proceed to the next step.
14. Press T to start the TFTP transfer.

    The firmware image is uploaded from the TFTP server and installed on the FIM which then reboots. When it starts up the module's configuration is reset to factory defaults. The module's configuration is synchronized to match the configuration of the primary module. The new module reboots again and can start processing traffic.
15. Verify that the configuration has been synchronized.

    The following command output shows the sync status of the FIMs in a FortiGate-7000 chassis. The field `in_sync=1` indicates that the configurations of the modules are synchronized.

    ```
    diagnose sys confsync
    status | grep in_sy
    FIM04E3E16000080, Slave, uptime=177426.45, priority=2,
    slot_id=1:2, idx=0, flag=0x0, in_sync=1
    FIM10E3E16000063, Master, uptime=177415.38, priority=1,
    slot_id=1:1, idx=1, flag=0x0, in_sync=1
    ```

    If `in_sync` is not equal to 1 or if a module is missing in the command output you can try restarting the modules in the chassis by entering execute reboot from any module CLI. If this does not solve the problem, contact Fortinet support.

Uploading firmware from a TFTP server to an FPM

Use the following steps to upload firmware from a TFTP server to an FPM. This procedure requires Ethernet connectivity between the TFTP server and one of the MGMT interfaces of one of the FIMs in the same FortiGate-7000 as the FPM.

During this procedure, the FPM will not be able to process traffic so, if possible, perform this procedure when the network is not processing any traffic. However, the other FPMs and the FIMs in the chassis should continue to operate normally and the chassis can continue processing traffic.

If you are operating an HA configuration, you should remove the FortiGate-7000 from the HA configuration before performing this procedure.

1. Set up a TFTP server and copy the firmware file to be installed into the TFTP server default folder.
2. Set up your network to allow traffic between the TFTP server and a MGMT interface of one of the FIMs in the chassis that also includes the FPM.

   You can use any MGMT interface of either of the FIMs. If the MGMT interface you are using is one of the MGMT interfaces connected as a LAG to a switch you must shutdown or disconnect all of the other connections in the LAG from the switch. This includes the MGMT interfaces in the other FIM.
3. Connect the console cable supplied with your chassis to the Console 1 port on your chassis front panel and to your management computer's RS-232 console port.

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. Press Ctrl-T to enter console switch mode.
6. Repeat pressing Ctrl-T until you have connected to the module to be updated. Example prompt:
`<Switching to Console: FPM03 (9600)>`
7. Optionally log into the FPM's CLI.
8. Reboot the FPM to be updated.
You can do this using the `execute reboot` command from the CLI or by pressing the power switch on the module front panel.
9. When the FPM starts up, follow the boot process in the terminal session and press any key when prompted to interrupt the boot process.
10. Press `C` to set up the TFTP configuration.
11. Use the BIOS menu to set the following. Only change settings if required.
`[P]: Set image download port`: The name of the FIM that can connect to the TFTP server (FIM01 is the FIM in slot 1 and FIM02 is the FIM in slot 2).
`[D]: Set DHCP mode`: Disabled.
`[I]: Set local IP address`: A temporary IP address to be used to connect to the TFTP server. This address must not be the same as the chassis 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`: The VLAN ID of the FIM interface that can connect to the TFTP server:

FIM01 local VLAN IDs

Interface	MGMT1	MGMT2	MGMT3	MGMT4
Local VLAN ID	11	12	13	14

FIM02 local VLAN IDs

Interface	MGMT1	MGMT2	MGMT3	MGMT4
Local VLAN ID	21	22	23	24

`[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 file to be installed.
12. Press Q to quit this menu.
13. Press R to review the configuration.
If you need to make any corrections, press C and make the changes as required. When the configuration is correct proceed to the next step.
14. Press T to start the TFTP transfer.
The firmware image is uploaded from the TFTP server and installed on the FPM which then reboots. When it starts up the module's configuration is reset to factory defaults. The module's configuration is synchronized to match the configuration of the primary module. The new module reboots again and can start processing traffic.
15. Verify that the configuration has been synchronized.
The following command output shows the sync status of the FIMs in a FortiGate-7000 chassis. The field `in_sync=1` indicates that the configurations of the modules are synchronized.

```
diagnose sys confsync
status | grep in_sy
```

```
FIM04E3E16000080, Slave, uptime=177426.45, priority=2,
slot_id=1:2, idx=0, flag=0x0, in_sync=1
FIM10E3E16000063, Master, uptime=177415.38, priority=1,
slot_id=1:1, idx=1, flag=0x0, in_sync=1
```

If `in_sync` is not equal to 1 or if a module is missing in the command output you can try restarting the modules in the chassis by entering execute reboot from any module CLI. If this does not solve the problem, contact Fortinet support.

Viewing interface statistics

From the primary FIM, you can add Interface History dashboard widgets to view traffic in and traffic out and total traffic information about the traffic passing through any FortiGate-7000 interface. You can add an Interface Bandwidth widget to the dashboard for:

- Any physical interface
- Any link aggregation (LAG) interface
- Any member of a LAG
- Any VLAN interface
- Any IPsec VPN tunnel interface
- Any redundant interface

The displayed data includes all of the traffic processed by the interface, independent of how the traffic is load balanced.

You can add individual Interface History widgets for each interface that you want to monitor. When adding the widget you can select the interface to monitor and select different time periods for which to display data. The data displayed on the widget updates in real time if you select enable refresh. You can also scan your cursor over the graph on the widget to see traffic in and out of the interface, and the total traffic for different times on the graph.

From the CLI, for physical interfaces, you can use the `diagnose hardware deviceinfo nic <interface-name>` command to view transmitted and received packets.

For LAG and VLAN interfaces you can display transmitted and received packets using the `diagnose netlink interface list <interface-name>` command.

FortiGate-7000 IPsec VPN

The following notes and limitations apply to FortiGate-7000 IPsec VPNs for FortiOS 5.4:

- Site-to-Site IPsec VPN is supported.
- Interface-based IPsec VPN (also called route-based IPsec VPN) is supported. Policy-based IPsec VPN is not supported.
- IPsec tunnels are not load-balanced across the FPMs, all IPsec tunnel sessions are sent to the primary (master) FPM.
- Traffic between IPsec VPN tunnels is supported.
- Phase 2 selectors are required to specify the IP addresses and netmasks of the source and destination subnets of the VPN. For more information, see Adding source and destination subnets to IPsec VPN phase 2 configurations on page 57.
- Remote networks with 0- to 15-bit netmasks are not supported. Remote networks with 16- to 32-bit netmasks are supported.
- IPv6 clear-text traffic (IPv6 over IPv4 or IPv6 over IPv6) is not supported.
- Dialup IPsec VPN is supported. The FortiGate-7000 can be the dialup server or client.
- Dynamic routing and policy routing is not supported for IPsec interfaces.
- IPsec static routes don't consider distance, weight, priority settings. IPsec static routes are always installed in the routing table, regardless of the tunnel state.
- Auto-negotiate is not supported for IPsec VPN phase 2 configurations. IPsec VPN tunnels should not be started by the FortiGate-7000 because this could lead to every FPM attempting to start an IPsec VPN tunnel with the same remote device.
- IPsec SA synchronization between HA peers is not supported. After an HA failover, IPsec VPN tunnels have to be re-initialized.

Adding source and destination subnets to IPsec VPN phase 2 configurations

FortiGate-7000 IPsec VPNs require phase 2 selectors. The phase 2 selectors specify the IP addresses and netmasks of the source and destination subnets of the VPN. The phase 2 selectors are mandatory on the FortiGate-7000 and are used to make sure that all IPsec VPN traffic is sent to the primary (master) FPM.

Use the following command to add phase 2 selectors.

```
config vpn ipsec phase2-interface
   edit "to_fgt2"
      set phase1name <name>
      set src-subnet <IP> <netmask>
      set dst-subnet <IP> <netmask>
   end
```

Where

`src-subnet` the subnet protected by the FortiGate that you are configuring and from which users connect to the destination subnet.

`dst-subnet` the destination subnet behind the remote IPsec VPN endpoint.

Example basic IPsec VPN phase 2 configuration

In a simple configuration such as the one below with an IPsec VPN between two remote subnets you can add the phase 2 selectors by adding the subnets to the phase 2 configuration as shown.

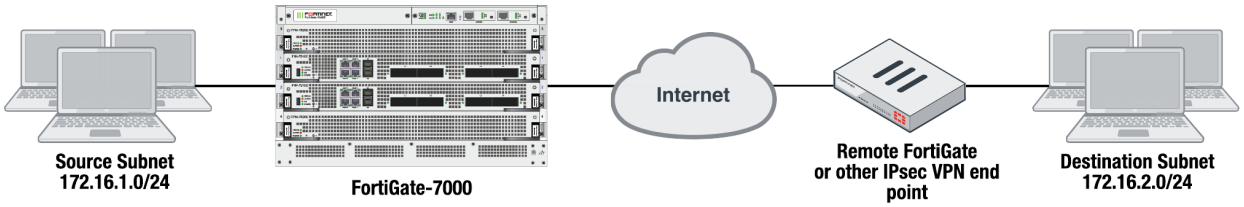

Enter the following command to add the source and destination subnets phase 2 selectors to the FortiGate-7000 IPsec VPN Phase 2 configuration.

```
config vpn ipsec phase2-interface
   edit "to_fgt2"So
      set phase1name "to_fgt2"
      set src-subnet 172.16.1.0 255.255.255.0
      set dst-subnet 172.16.2.0 255.255.255.0
   end
```

Example multiple subnet IPsec VPN phase 2 configuration

In a more complex configuration, such as the one below with a total of 5 subnets you still need to use the phase 2 selectors to add all of the subnets to the Phase 2 configuration. In this case you can create a firewall address for each subnet, add the addresses to address groups, and add the address groups to the phase 2 selectors.

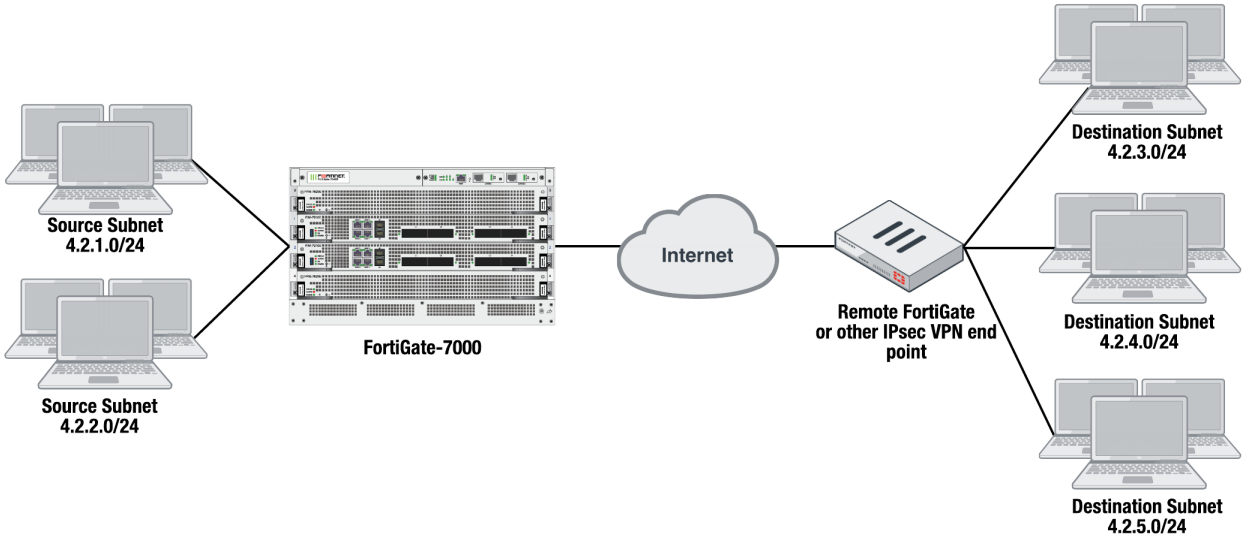

Enter the following commands to create firewall addresses for each subnet.

```
config firewall address
   edit "local_subnet_1"
      set subnet 4.2.1.0 255.255.255.0
   next
   edit "local_subnet_2"
      set subnet 4.2.2.0 255.255.255.0
   next
   edit "remote_subnet_3"
      set subnet 4.2.3.0 255.255.255.0
   next
   edit "remote_subnet_4"
      set subnet 4.2.4.0 255.255.255.0
   next
   edit "remote_subnet_5"
      set subnet 4.2.5.0 255.255.255.0
   end
```

And then put the five firewall addresses into two firewall address groups.

```
config firewall addrgrp
   edit "local_group"
      set member "local_subnet_1" "local_subnet_2"
   next
   edit "remote_group"
      set member "remote_subnet_3" "remote_subnet_4" "remote_subnet_5"
   end
```

Now, use the firewall address groups in the Phase 2 configuration:

```
config vpn ipsec phase2-interface
   edit "to-fgt2"
      set phase1name "to-fgt2"
      set src-addr-type name
      set dst-addr-type name
      set src-name "local_group"
```

```
    set dst-name "remote_group"
  end
```

Configuring the FortiGate-7000 as a dialup IPsec VPN server

FortiGate-7000s running v5.4.5 can be configured as dialup IPsec VPN servers.

Example dialup IPsec VPN configuration

The following shows how to setup a dialup IPsec VPN configuration where the FortiGate-7000 acts as a dialup IPsec VPN server.

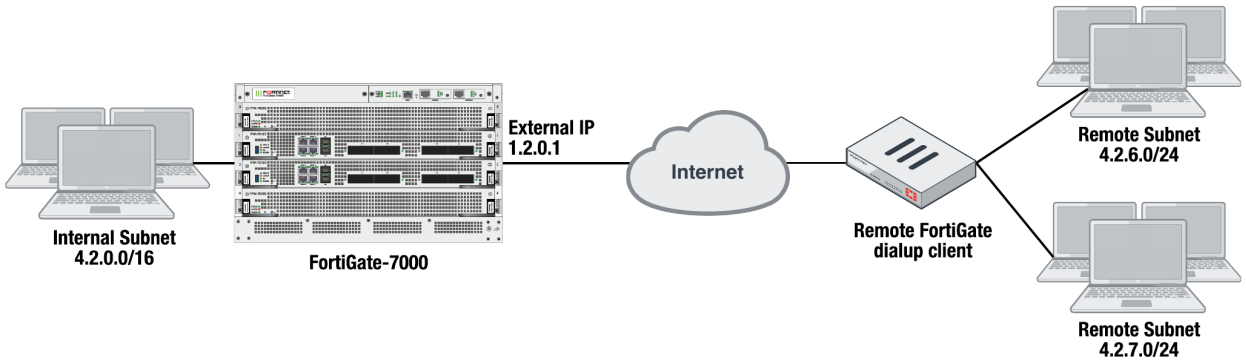

To configure the FortiGate-7000 as a dialup IPsec VPN server:

Configure the phase1, set type to dynamic.

```
config vpn ipsec phase1-interface
  edit dialup-server
    set type dynamic
    set interface "v0020"
    set peertype any
    set psksecret < password>
  end
```

Configure the phase 2, to support dialup IPsec VPN, set the destination subnet to 0.0.0.0 0.0.0.0.

```
config vpn ipsec phase2-interface
  edit dialup-server
    set phase1name dialup-server
    set src-subnet 4.2.0.0 255.255.0.0
    set dst-subnet 0.0.0.0 0.0.0.0
  end
```

To configure the remote FortiGate as a dialup IPsec VPN client

The dialup IPsec VPN client should advertise its local subnet(s) using the phase 2 src-subnet option.

If there are multiple local subnets create a phase 2 for each one. Each phase 2 only advertises one local subnet to the dialup IPsec VPN server. If more than one local subnet is added to the phase 2, only the first one is advertised to the server.

Dialup client configuration:

```
config vpn ipsec phase1-interface
   edit "to-fgt7k"
      set interface "v0020"
      set peertype any
      set remote-gw 1.2.0.1
      set psksecret <password>
   end
   config vpn ipsec phase2-interface
   edit "to-fgt7k"
      set phase1name "to-fgt7k"
      set src-subnet 4.2.6.0 255.255.255.0
      set dst-subnet 4.2.0.0 255.255.0.0
   next
   edit "to-fgt7k-2"
      set phase1name "to-fgt7k"
      set src-subnet 4.2.7.0 255.255.255.0
      set dst-subnet 4.2.0.0 255.255.0.0
   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 primary FIM interface module to determine the primary FPM processor module. For FortiGate-7000 HA, run this command from the primary FortiGate-7000. The third line of the command output shows which FPM is operating as the primary FPM.

```
diagnose load-balance status
  FIM01: FIM04E3E16000074
  Master FPM Blade: slot-4

     Slot  3: FPM20E3E17900113
       Status:Working   Function:Active
       Link:      Base: Up          Fabric: Up
       Heartbeat: Management: Good   Data: Good
       Status Message:"Running"
     Slot  4: FPM20E3E16800033
       Status:Working   Function:Active
       Link:      Base: Up          Fabric: Up
       Heartbeat: Management: Good   Data: Good
       Status Message:"Running"
```

```
  FIM02: FIM10E3E16000040
  Master FPM Blade: slot-4

     Slot  3: FPM20E3E17900113
       Status:Working   Function:Active
       Link:      Base: Up          Fabric: Up
       Heartbeat: Management: Good   Data: Good
       Status Message:"Running"
     Slot  4: FPM20E3E16800033
       Status:Working   Function:Active
       Link:      Base: Up          Fabric: Up
       Heartbeat: Management: Good   Data: Good
       Status Message:"Running"
```

Log into the primary FPM CLI and from here log into the VDOM that you added the tunnel configuration to 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.

```
CH15 [FPM04] (002ipsecvpn) # 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 FIMs 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
       ========================================
```

High availability

FortiGate-7000 supports a variation of active-passive FortiGate Clustering Protocol (FGCP) high availability between two identical FortiGate-7000s. With active-passive FortiGate-7000 HA, you create redundant network connections to two identical FortiGate-7000s and add redundant HA heartbeat connections. Then you configure the FIMs in each FortiGate-7000 for HA.The FGCP forms a cluster and selects a primary FortiGate-7000.

Example FortiGate-7040 HA configuration

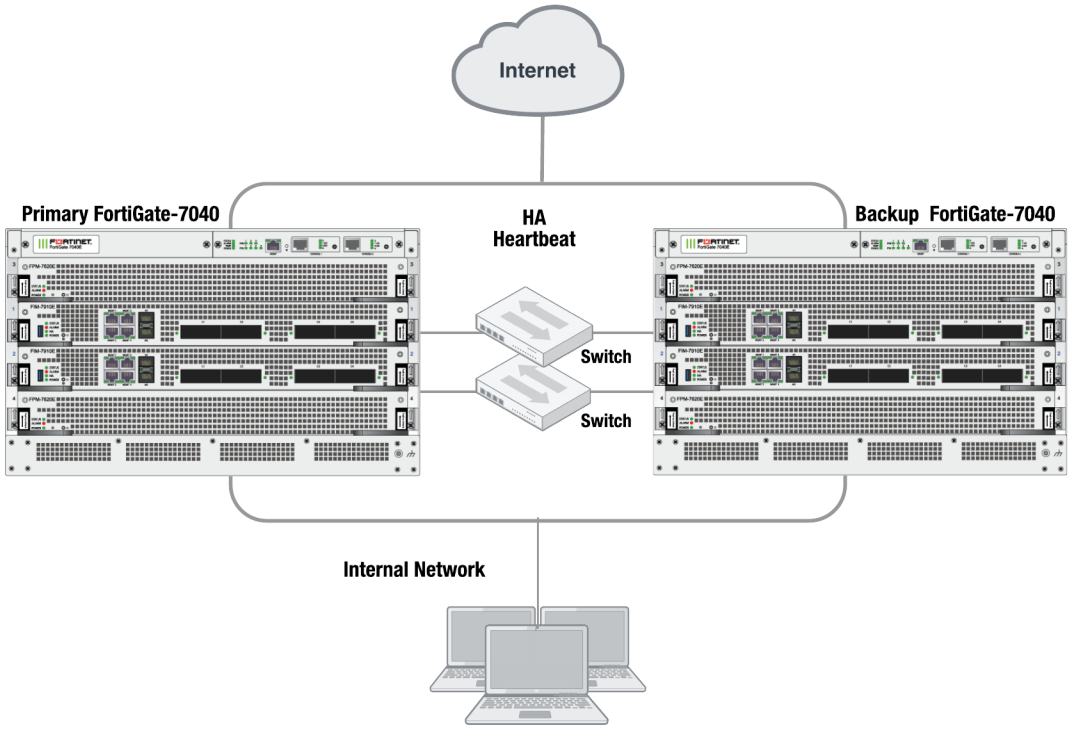

The primary FortiGate-7000 processes all traffic. The backup FortiGate-7000 operates in hot standby mode. The FGCP synchronizes the configuration, active sessions, routing information, and so on to the backup FortiGate-7000. If the primary FortiGate-7000 fails, traffic automatically fails over to the backup FortiGate-7000.

The FGCP selects the primary FortiGate-7000 based on a number of criteria including the configured priority, the bandwidth, the number of FIM interface failures, and the number of FPM or FIMs that have failed. As part of the HA configuration you assign each FortiGate-7000 a chassis ID and you can set the priority of each FIM and configure module failure tolerances and link failure thresholds.

Before you begin configuring HA

Before you begin, the FortiGate-7000s should be running the same FortiOS firmware version and interfaces should not be configured to get their addresses from DHCP or PPPoE. Register and apply licenses to the each FortiGate-7000 before setting up the HA cluster. This includes licensing for FortiCare, IPS, AntiVirus, Web Filtering, Mobile Malware, FortiClient, FortiCloud, and additional virtual domains (VDOMs). Both FortiGate-7000s in the cluster must have the same level of licensing for FortiGuard, FortiCloud, FortiClient, and VDOMs. FortiToken licenses can be added at any time because they are synchronized to all cluster members.

If required, you should configure split ports on the FIMs on both chassis before configuring HA because the modules have to reboot to enable the split port configuration. For example, to split the C1, C2, and C4 interfaces of an FIM-7910E in slot 1, enter the following command:

```
config system global
   set split-port 1-C1 2-C1 2-C4
end
```

After configuring split ports, the FortiGate-7000 reboots and synchronizes the configuration.

On each FortiGate-7000, make sure configurations of the modules are synchronized before starting to configure HA. You can use the following command to verify that the configurations of all of the modules are synchronized:

```
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
all: c0 68 d2 67 e1 23 d9 3a 10 50 45 c5 50 f1 e6 8e
```

If the modules are synchronized, the checksums displayed should all be the same.

You can also use the following command to list the modules that are synchronized. The example output shows all four modules in a FortiGate-7040E have been configured for HA and added to the cluster.

```
diagnose sys configsync status | grep in_sync
Master, uptime=692224.19, priority=1, slot_1d=1:1, idx=0, flag=0x0, in_sync=1
Slave, uptime=676789.70, priority=2, slot_1d=1:2, idx=1, flag=0x0, in_sync=1
Slave, uptime=692222.01, priority=17, slot_1d=1:4, idx=2, flag=0x64, in_sync=1
Slave, uptime=692271.30, priority=16, slot_1d=1:3, idx=3, flag=0x64, in_sync=1
```

In this command output `in_sync=1` means the module is synchronized with the primary FIM and `in_sync=0` means the module is not synchronized.

Connect the M1 and M2 interfaces for HA heartbeat communication

HA heartbeat communication between FortiGate-7000s happens over the 10Gbit M1 and M2 interfaces of the FIMs in each chassis. To set up HA heartbeat connections:

- Connect the M1 interfaces of all FIMs together using a switch.
- Connect the M2 interfaces of all FIMs together using another switch.

All of the M1 interfaces must be connected together with a switch and all of the M2 interfaces must be connected together with another switch. Connecting M1 interfaces or M2 interfaces directly is not supported as each FIM needs to communicate with all other FIMs.

Connect the M1 and M2 interfaces before enabling HA. Enabling HA moves heartbeat communication between the FIMs in the same chassis to the M1 and M2 interfaces. So if these interfaces are not connected before you enable HA, FIMs in the same chassis will not be able to communicate with each other.

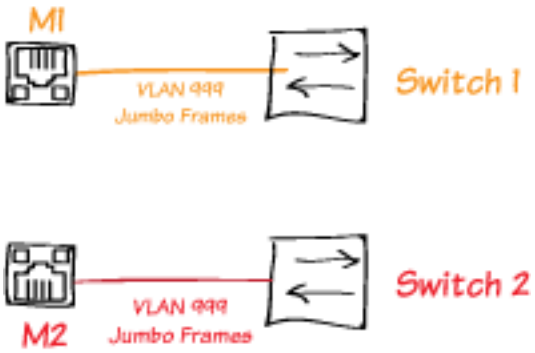

Heartbeat packets are VLAN packets with VLAN ID 999 and ethertype 9890. The MTU value for the M1 and M2 interfaces is 1500.

You can use the following command to change the HA heartbeat packet VLAN ID and ethertype values if required for your switches. You must change these settings on each FIM interface module. By default the M1 and M2 interface heartbeat packets use the same VLAN IDs. The following example changes the M1 VLAN ID to 991 and the M2 VLAN ID to 992.

```
config system ha
   set hbdev M1/M2
   set hbdev-vlan-id 991
   set hbdev-second-vlan-id 992
   set ha-eth-type <eth-type>
end
```

For this configuration to work, you must configure both VLAN IDs. You cannot use the default value of 999.

Recommended HA heartbeat interface configuration

For redundancy, Fortinet recommends using separate switches for the M1 and M2 connections.These switches should be dedicated to HA heartbeat communication and not used for other traffic.

If you use the same switch for the M1 and M2 interfaces, separate the M1 and M2 traffic on the switch and set the heartbeat traffic on the M1 and M2 interfaces to have different VLAN IDs.

If you don't set different VLAN IDs for the M1 and M2 heartbeat packets, you must enable q-in-q on the switch.

Sample switch configuration

Sample switch configuration for a Cisco Catalyst switch. This configuration sets the interface speeds, configures the switch to allow vlan 999, and enables trunk mode:

```
##interface config
interface TenGigabitEthernet1/0/5
description Chassis1 FIM1 M1
switchport trunk allowed vlan 999
switchport mode trunk
```

If you are using one switch for both M1 and M2 connections, the configuration would be the same except you would add q-in-q support and two different VLANs, one for M1 traffic and one for M2 traffic.

For the M1 connections:

```
interface Ethernet1/5
description QinQ Test
switchportmode dot1q-tunnel
switchport access vlan 777
spanning-tree port type edge
```

For the M2 connections:

```
interface Ethernet1/5
description QinQ Test
switchport mode dot1q-tunnel
switchport access vlan 888
spanning-tree port type edge
```

HA packets must have the configured VLAN tag. If the switch removes or changes this tag, HA heartbeat communication will not work and network traffic will be disrupted.

HA configuration

Use the following steps to setup the configuration for HA between two FortiGate-7000s (in this example two FortiGate-7040E or 7060Es). Similar steps are required for the FortiGate-7030E except that each FortiGate-7030E only has one FIM.

Each FIM has to be configured for HA separately. The HA configuration is not synchronized among FIMs. You can begin by setting up the first FortiGate-7000 and setting up HA on both of the FIMs in it. Then do the same for the second FortiGate-7000.

To configure HA you assign a chassis ID (1 and 2) to each of the FortiGate-7000s is assigned a chassis ID.These IDs just allow you to identify the chassis and do not influence primary unit selection.

Setting up HA on the FIMs in the first FortiGate-7000 (chassis 1)

1. Log into the CLI of the FIM in slot 1 (FM01) and enter the following command:

```
config system ha
   set mode a-p
   set password <password>
   set group-id <id>
```

```
   set chassis-id 1
   set hbdev M1/M2
end
```

This adds basic HA settings to this FIM.

2. Repeat this configuration on the FIM in slot 2 (FIM02).

```
config system ha
   set mode a-p
   set password <password>
   set group-id <id>
   set chassis-id 1
   set hbdev M1/M2
end
```

3. From either FIM, enter the following command to confirm that the FortiGate-7000 is in HA mode:

```
diagnose sys ha status
```

The `password` and `group-id` are unique for each HA cluster and must be the same on all FIMs. If a cluster does not form, one of the first things to check are `groupd-id` and re-enter the password on both FIMs.

Configure HA on the FIMs in the second FortiGate-7000 (chassis 2)

1. Repeat the same HA configuration settings on the FIMs in the second chassis except set the chassis ID to 2.

```
config system ha
   set mode a-p
   set password <password>
   set group-id <id>
   set chassis-id 2
   set hbdev M1/M2
end
```

2. From any FIM, enter the following command to confirm that the cluster has formed and all of the FIMs have been added to it:

```
diagnose sys ha status
```

The cluster has now formed and you can add the configuration and connect network equipment and start operating the cluster. You can also modify the HA configuration depending on your requirements.

Verifying that the cluster is operating correctly

Enter the following CLI command to view the status of the cluster. You can enter this command from any module's CLI. The HA members can be in a different order depending on the module CLI from which you enter the command.

If the cluster is operating properly the following command output should indicate the primary and backup (master and slave) FortiGate-7000 as well as primary and backup (master and slave) modules. For each module, the `state` portion of the output shows all the parameters used to select the primary FIM. These parameters include the number FPMs that the FIM is connecting to that have failed, the status of any link aggregation group (LAG) interfaces in the configuration, the state of the interfaces in the FIM, the traffic bandwidth score for the FIM (the higher the traffic bandwidth score the more interfaces are connected to networks), and the status of the management links.

```
diagnose  sys  ha  status
==========================================================================
Current slot: 1  Module SN: FIM04E3E16000085
Chassis HA mode: a-p

Chassis HA information:
```

```
[Debug_Zone HA information]
HA group member information: is_manage_master=1.
FG74E83E16000015:  Slave, serialno_prio=1, usr_priority=128, hostname=CH15
FG74E83E16000016: Master, serialno_prio=0, usr_priority=127, hostname=CH16

HA member information:
CH16(FIM04E3E16000085), Master(priority=0), uptime=78379.78, slot=1, chassis=2(2)
    slot: 1, chassis_uptime=145358.97,
    more: cluster_id:0, flag:1, local_priority:0, usr_priority:127, usr_override:0
    state: worker_failure=0/2, lag=(total/good/down/bad-score)=5/5/0/0,
           intf_state=(port up)=0, force-state(1:force-to-master)
           traffic-bandwidth-score=120, mgmt-link=1
    hbdevs: local_interface=      1-M1 best=yes
            local_interface=      1-M2 best=no
    ha_elbc_master: 3, local_elbc_master: 3

CH15(FIM04E3E16000074), Slave(priority=2), uptime=145363.64, slot=1, chassis=1(2)
    slot: 1, chassis_uptime=145363.64,
    more: cluster_id:0, flag:0, local_priority:2, usr_priority:128, usr_override:0
    state: worker_failure=0/2, lag=(total/good/down/bad-score)=5/5/0/0,
           intf_state=(port up)=0, force-state(-1:force-to-slave)
           traffic-bandwidth-score=120, mgmt-link=1
    hbdevs: local_interface=      1-M1 last_hb_time=145640.39   status=alive
            local_interface=      1-M2 last_hb_time=145640.39   status=alive

CH15(FIM10E3E16000040), Slave(priority=3), uptime=145411.85, slot=2, chassis=1(2)
    slot: 2, chassis_uptime=145638.51,
    more: cluster_id:0, flag:0, local_priority:3, usr_priority:128, usr_override:0
    state: worker_failure=0/2, lag=(total/good/down/bad-score)=5/5/0/0,
           intf_state=(port up)=0, force-state(-1:force-to-slave)
           traffic-bandwidth-score=100, mgmt-link=1
    hbdevs: local_interface=      1-M1 last_hb_time=145640.62   status=alive
            local_interface=      1-M2 last_hb_time=145640.62   status=alive

CH16(FIM10E3E16000062), Slave(priority=1), uptime=76507.11, slot=2, chassis=2(2)
    slot: 2, chassis_uptime=145641.75,
    more: cluster_id:0, flag:0, local_priority:1, usr_priority:127, usr_override:0
    state: worker_failure=0/2, lag=(total/good/down/bad-score)=5/5/0/0,
           intf_state=(port up)=0, force-state(-1:force-to-slave)
           traffic-bandwidth-score=100, mgmt-link=1
    hbdevs: local_interface=      1-M1 last_hb_time=145640.39   status=alive
            local_interface=      1-M2 last_hb_time=145640.39   status=alive
```

Setting up HA management connections

Fortinet recommends the following configurations for redundant management connections to a FortiGate-7000 HA configuration.

- Single management connections to each of the FIMs.
- Redundant management connections to each of the FIMs.

These management connections involve connecting the static redundant management interfaces (MGMT1 to MGMT4) of each FIM in the HA configuration to one or more switches. You do not have to change the FortiGate-7000

configuration to set up redundant management connections. However, specific switch configurations are required for each of these configurations as described below.

LACP is not supported for the mgmt aggregate interface.

Setting up single management connections to each of the FIMs

The simplest way to provide redundant management connections to a FortiGate-7000 HA configuration involves connecting the MGMT1 interface of each of the FIMs to four ports on a switch. On the switch you must add the four switch ports to the same VLAN. Then connect the switch to your management network and allow traffic from the VLAN to the management network.

A FortiGate-7030E HA configuration only has two FIMs so would only require two switch ports.

Example FortiGate-7000 HA redundant management connections

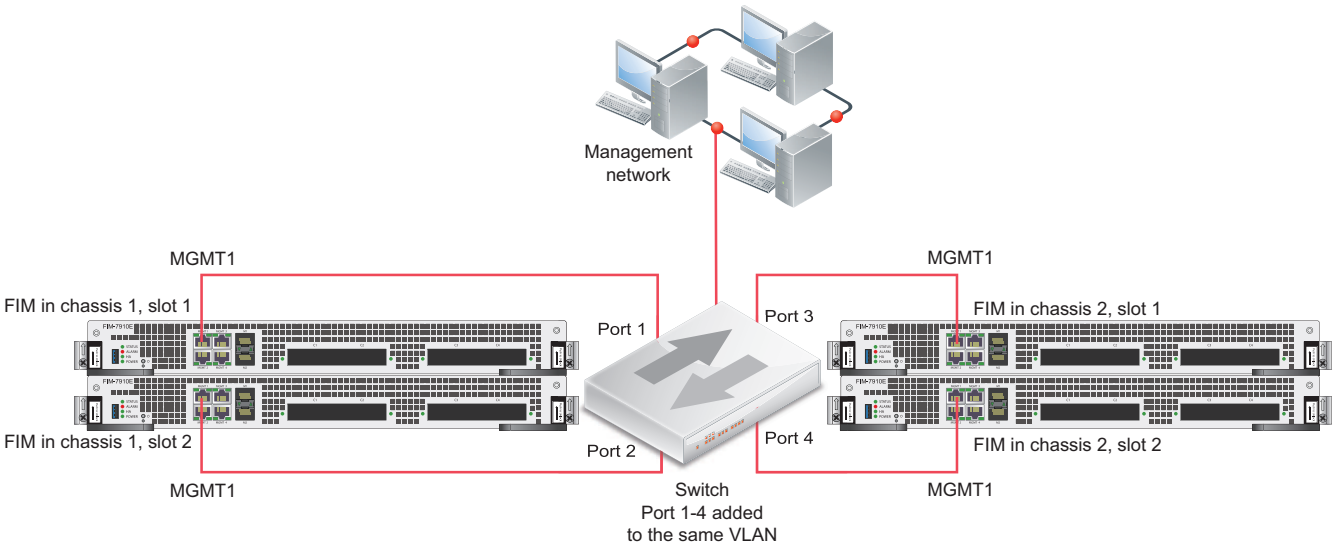

Setting up redundant management connections to each of the FIMs

You can enhance redundancy by setting up two redundant management connections to each FIM. To support this configuration, on the switch you must create a port channel for each FIM interface. Create a total of four port channels, one for each FIM and add each of the port channels to the same VLAN. Then connect the switch to your management network and allow traffic from the VLAN to the management network.

If you use two switches, the VLAN should span across both switches.

A FortiGate-7030E HA configuration only has two FIMs so would only require two port channels.

Example FortiGate-7000 HA redundant management connections with redundant connections to each FIM

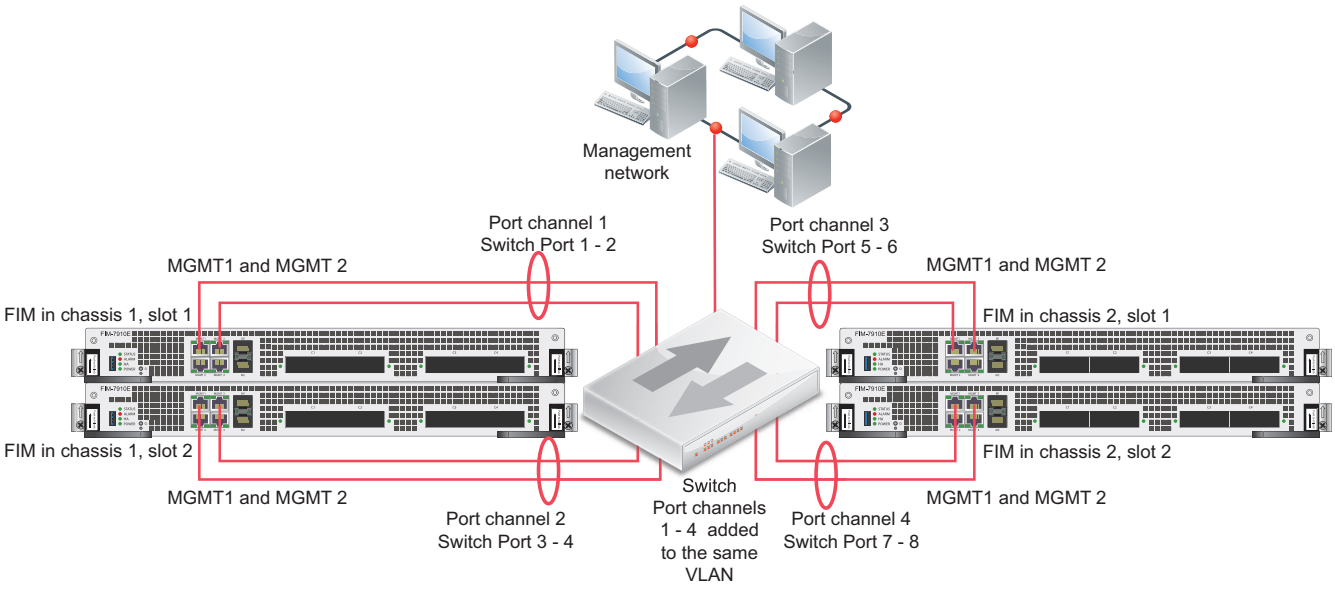

Managing individual modules in HA mode

In some cases you may want to connect to an individual FIM or FPM in a specific FortiGate-7000. For example, you may want to view the traffic being processed by the FPM in slot 3 of the FortiGate-7000 designated as chassis 2. You can connect to the GUI or CLI of individual modules in the chassis using the system management IP address with a special port number.

For example, if the system management IP address is 1.1.1.1 you can browse to https://1.1.1.1:44323 to connect to the FPM in the FortiGate-7000 chassis 2 slot 3. The special port number (in this case 44323) is a combination of the service port, chassis ID, and slot number. The following table lists the special ports for common admin protocols:

FortiGate-7000 HA special administration port numbers

Chassis and Slot Number	Slot Address	HTTP (80)	HTTPS (443)	Telnet (23)	SSH (22)	SNMP (161)
Ch1 slot 5	FPM05	8005	44305	2305	2205	16105
Ch1 slot 3	FPM03	8005	44303	2303	2203	16103
Ch1 slot 1	FIM01	8003	44301	2301	2201	16101
Ch1 slot 2	FIM02	8002	44302	2302	2202	16102

Chassis and Slot Number	Slot Address	HTTP (80)	HTTPS (443)	Telnet (23)	SSH (22)	SNMP (161)
Ch1 slot 4	FPM04	8004	44304	2304	2204	16104
Ch1 slot 6	FPM06	8006	44306	2306	2206	16106
Ch2 slot 5	FPM05	8005	44325	2325	2225	16125
Ch2 slot 3	FPM03	8005	44323	2323	2223	16123
Ch2 slot 1	FIM01	8003	44321	2321	2221	16121
Ch2 slot 2	FIM02	8002	44322	2322	2222	16122
Ch2 slot 4	FPM04	8004	44324	2324	2224	16124
Ch2 slot 6	FPM06	8006	44326	2326	2226	16126

For example:

- To connect to the GUI of the FPM in chassis 1 slot 3 using HTTPS you would browse to https://1.1.1.1:44313.
- To send an SNMP query to the FPM in chassis 2 slot 6 use the port number 16126.

The formula for calculating the special port number is based on Chassis ID. CH1 = Chassis ID1, CH2 = Chassis ID2. The formula is: service_port x 100 + (chassis_id – 1) x 20 + slot_id.

Upgrading FortiGate-7000 HA cluster firmware

All of the modules in your FortiGate-7000 HA configuration run the same firmware image. You upgrade the firmware from the primary FIM GUI or CLI just as you would any FortiGate product. During the upgrade process, the firmware of all of the modules in both chassis upgrades automatically in the background. Firmware upgrades should be done during a quiet time because traffic may be briefly be interrupted during the upgrade process.

Before beginning a firmware upgrade, Fortinet recommends the following:

- 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.
- Backup your FortiGate-7000 HA configuration.
- Review the services provided by your FortiGate-7000 before the upgrade and then again after the upgrade to make sure everything continues to operate normally. For example, you might want to verify that you can successfully access a key server used by your organization before the upgrade and make sure after the upgrade that you can still reach the server and that performance is comparable. You could also take a snapshot of key performance indicators (number of sessions, CPU usage, and memory usage) before the upgrade and verify that you see comparable performance after the upgrade .

Enabling uninterruptable-upgrade

Fortinet recommends upgrading your FortiGate-7000 HA configuration firmware with `uninterruptable-upgrade` enabled. With `uninterruptable-upgrade` enabled, the FortiGate-7000 HA configuration goes through a multi-step process to upgrade the firmware of all components in the configuration. Since many components are involved, the entire upgrade process may take a few minutes. It all happens automatically and should cause only minor traffic disruptions. Because of the possible disruptions, you should upgrade HA cluster firmware when traffic is low or during a maintenance period.

Use the following command to enable uninterruptable-upgrade:

```
config system ha
   set uninterruptable-upgrade enable
end
```

The following steps happen in the background when upgrading the firmware running on a FortiGate-7000 HA cluster with `uninterruptable-upgrade` enabled.

- The firmware upgrade downloads to the primary (master) FortiGate-7000.
- The primary FortiGate-7000 sends a copy of the firmware upgrade file to the backup (slave) FortiGate-7000.
- The backup FortiGate-7000 upgrades its firmware, restarts, and rejoins the cluster.
- The primary FortiGate-7000 verifies that all members of the backup FortiGate-7000 can process traffic. The firmware upgrade will not proceed until all of the backup FortiGate-7000 components are operating.
- The primary FortiGate-7000 then sends a switchover command and the backup FortiGate-7000 becomes the primary FortiGate-7000.
- The new primary FortiGate-7000 sends gratuitous ARP packets to inform attached network devices to send packets to the new primary FortiGate-7000.
- Traffic switches over to the new primary FortiGate-7000.
- The original primary FortiGate-7000 upgrades its firmware, restarts, and rejoins the cluster as the backup FortiGate-7000.

The amount of time this process takes and the probability of minor traffic disruptions depends on the number of modules in your FortiGate-7000 and on traffic load conditions, the network configuration, and how quickly the gratuitous ARP packets update network devices.

Distributed clustering

FortiGate-7000 HA supports separating the FortiGate-7000s in different physical locations. Distributed FortiGate-7000 HA clustering (or geographically distributed FortiGate-7000 HA or geo clustering) can involve two FortiGate-7000s 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-7000 HA configuration, distributed FortiGate-7000 HA requires heartbeat communication between the FortiGate-7000s over the M1 and M2 interfaces. In a distributed FortiGate-7000 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 communication over the same subnet.

The M1 and M2 interface traffic must be separated. You can do this by using separate channels for each interface or by configuring the M1 and M2 interfaces to use different VLANs.

Example FortiGate-7000 distributed clustering configuration

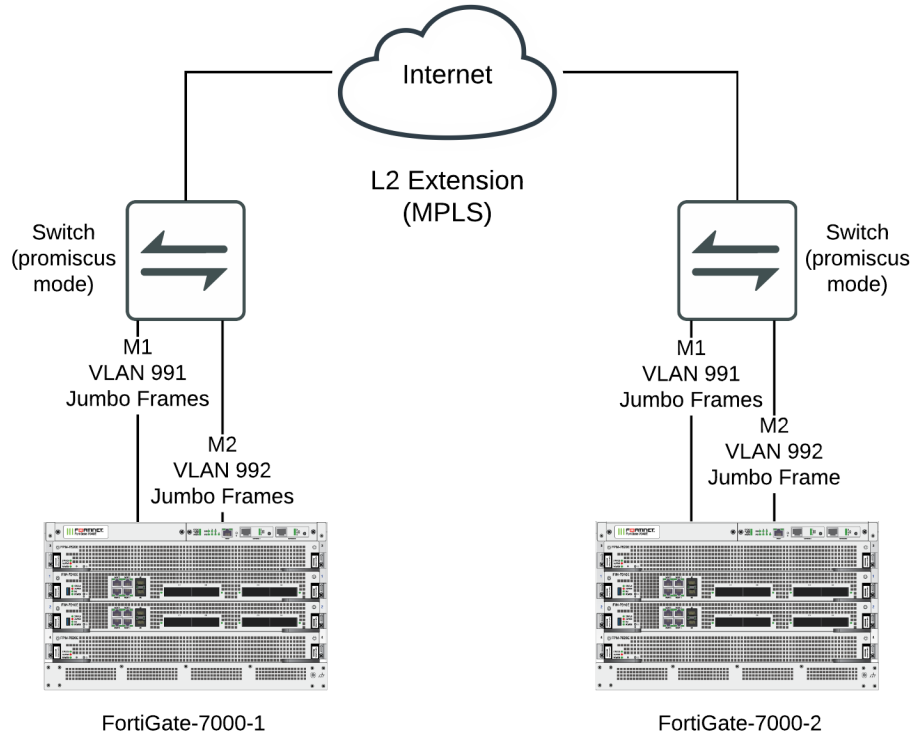

Because of the possible distance between sites, it may take a relatively long time for heartbeat packets to be transmitted between the FortiGate-7000s. This could lead to a split brain scenario. To avoid a split brain scenario you can modify heartbeat timing so that the cluster expects extra time between 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.

Modifying heartbeat timing

If the FortiGate-7000s in the HA cluster do not receive heartbeat packets on time, the FortiGate-7000s in the HA configuration may each determine that the other FortiGate-7000 has failed. HA heartbeat packets may not be sent on time because of network issues. For example, if the M1 and M2 communications links between the FortiGate-7000s become too busy to handle the heartbeat traffic. Also, in a distributed clustering configuration the round trip time (RTT) between the FortiGate-7000s may be longer the expected time between heartbeat packets.

In addition, if the FortiGate-7000s becomes excessively busy, they may delay sending heartbeat packets.

Even with these delays, the FortiGate-7000 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 does not receive before assuming that a failure has occurred. The default value of 6mean that if a FortiGate-7000 does not receive 6 heartbeat packets it determines that the other FortiGate-7000 in the cluster has failed. The range is 1 to 60 packets.

The lower the `hb-lost-threshold`, the faster a FortiGate-7000 HA configuration responds when a failure occurs. However, sometimes heartbeat packets may not be received because the other FortiGate-7000 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-7000 waits before assuming that the other FortiGate-7000 has failed and is no longer sending heartbeat packets. By default, if a FortiGate-7000 does not receive a heartbeat packet from a cluster unit for 6 * 200 = 1200 milliseconds or 1.2 seconds the FortiGate-7000 assumes that the other FortiGate-7000 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-7000 waits before changing from hello state to work state. After a failure or when starting up, FortiGate-7000s in HA mode operate in the hello state to send and receive heartbeat packets to find each other and form a cluster. A FortiGate-7000 should change from the hello state to work state after it finds the FortiGate-7000 to form a cluster with. If for some reason the FortiGate-7000s cannot find each other during the hello state both FortiGate-7000s may assume that the other one has failed and each could form separate clusters of one FortiGate-7000. The FortiGate-7000s 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-7000s finding each other could be the FortiGate-7000s are located at different sites or for some other reason communication is delayed between the heartbeat interfaces. If you find that your FortiGate-7000s 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
```

Session failover (session-pickup)

Session failover means that after a failover, communications sessions resume on the new primary FortiGate-7000 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

Session failover (also called session-pickup) is not enabled by default for FortiGate-7000 HA. If sessions pickup is enabled, while the FortiGate-7000 HA cluster is operating the primary FortiGate-7000 informs the backup FortiGate-7000 of changes to the primary FortiGate-7000 connection and state tables for TCP and UDP sessions passing through the cluster, keeping the backup FortiGate-7000 up-to-date with the traffic currently being processed by the cluster.

After a failover the new primary FortiGate-7000 recognizes open sessions that were being handled by the cluster. The sessions continue to be processed by the new primary FortiGate-7000 and are handled according to their last known state.

Session-pickup has some limitations. For example, session failover is not supported 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 VPN, 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 and UDP

To enable session-pickup, from the CLI enter:

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

When session-pickup is enabled, sessions in the primary FortiGate-7000 TCP and UDP session tables are synchronized to the backup FortiGate-7000. As soon as a new TCP or UDP session is added to the primary FortiGate-7000 session table, that session is synchronized to the backup FortiGate-7000. This synchronization happens as quickly as possible to keep the session tables synchronized.

If the primary FortiGate-7000 fails, the new primary FortiGate-7000 uses its synchronized session tables to resume all TCP and UDP sessions that were being processed by the former primary FortiGate-7000 with only minimal interruption. Under ideal conditions all TCP and UDP 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-7000 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-7000 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.

Primary FortiGate-7000 chassis selection and failover criteria

Once two FortiGate-7000s recognize that they can form a cluster, they negotiate to select a primary FortiGate-7000. Primary FortiGate-7000 selection occurs automatically based on the criteria shown below. After the cluster selects the primary FortiGate-7000, the other FortiGate-7000 becomes the backup.

Negotiation and primary FortiGate-7000 selection also takes place if the one of the criteria for selecting the primary chassis changes. For example, an interface can become disconnected or module can fail. After this happens, the cluster can renegotiate to select a new primary FortiGate-7000 also using the criteria shown below.

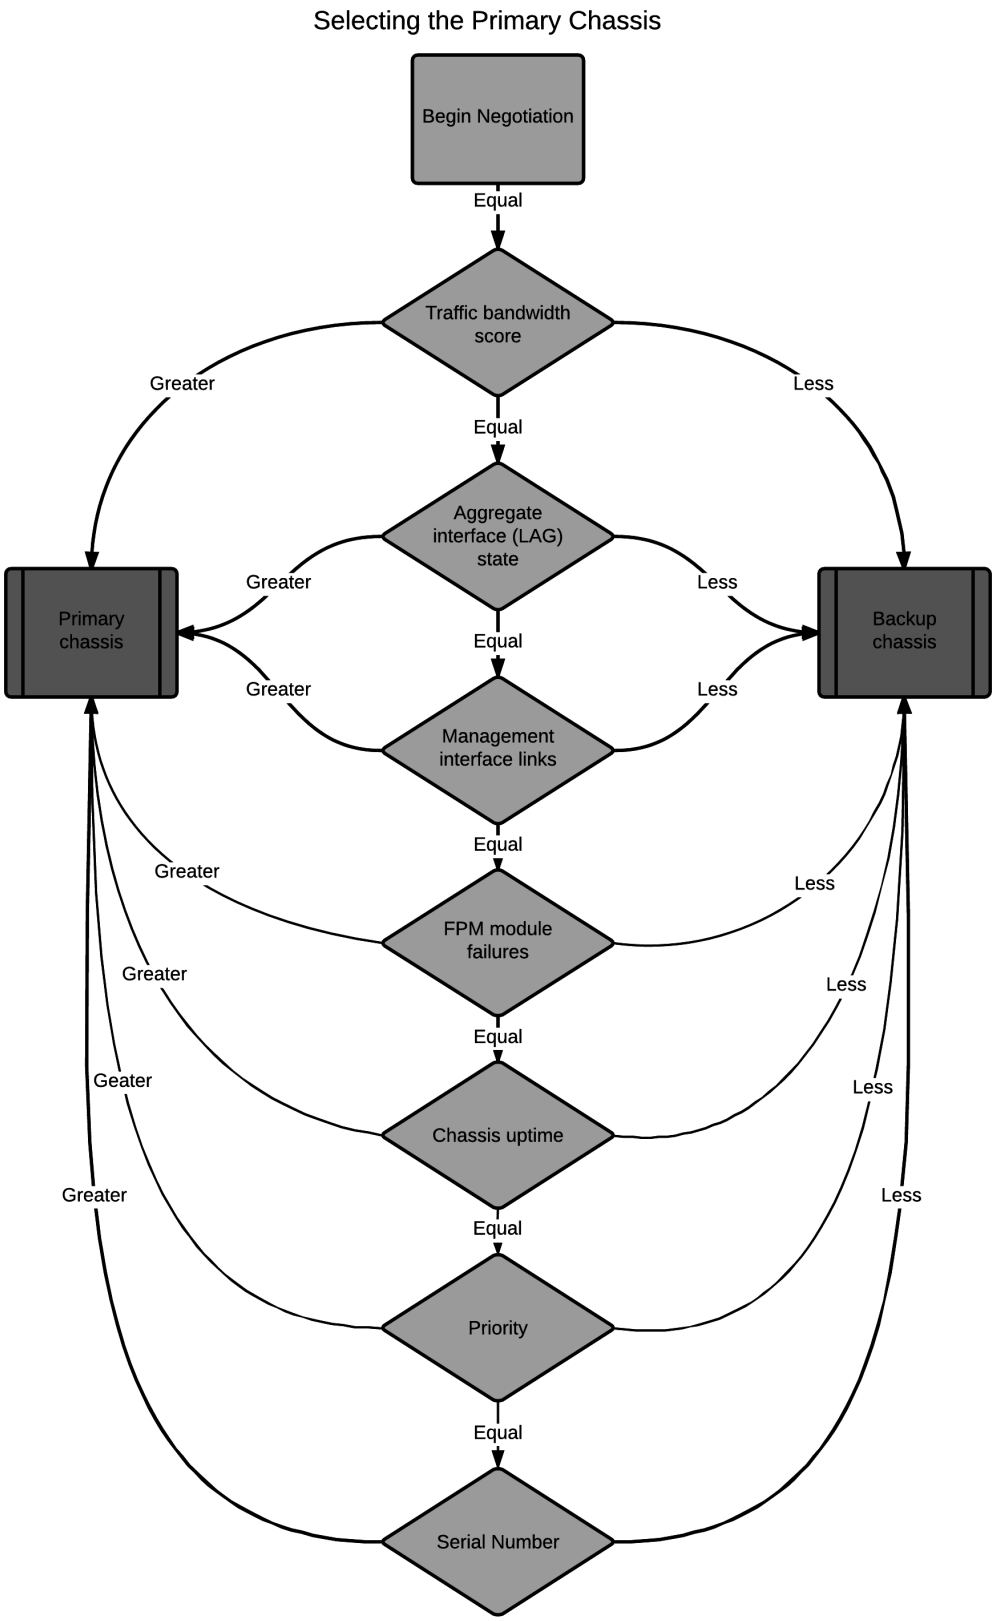
Selecting the Primary Chassis
Begin Negotiation
Equal
Traffic bandwidth score
Greater
Less
Equal
Aggregate interface (LAG) state
Greater
Less
Primary chassis
Backup chassis
Equal
Greater
Less
Management interface links
Equal
Greater
Less
FPM module failures
Greater
Less
Equal
Geater
Less
Chassis uptime
Greater
Less
Equal
Priority
Equal
Serial Number

If there are no failures and if you haven't configured any settings to influence primary FortiGate-7000 selection, the FortiGate-7000 with the highest serial number becomes the primary FortiGate-7000.

Using the serial number is a convenient way to differentiate FortiGate-7000s; so basing primary FortiGate-7000 selection on the serial number is predictable and easy to understand and interpret. Also the FortiGate-7000 with the highest serial number would usually be the newest with the most recent hardware version. In many cases you may not need active control over primary FortiGate-7000 selection, so basic primary FortiGate-7000 selection based on serial number is sufficient.

In some situations you may want have control over the FortiGate-7000 that becomes the primary FortiGate-7000. You can control primary FortiGate-7000 selection by setting the priority of one FortiGate-7000 to be higher than the priority of the other. If you change the priority of one of the FortiGate-7000, during negotiation, the FortiGate-7000 with the highest priority becomes the primary FortiGate-7000 . As shown above, FGCP selects the primary FortiGate-7000 based on priority before serial number. For more information about how to use priorities, see **Primary FortiGate-7000 chassis selection and failover criteria on page 78**.

FortiGate-7000 uptime is also a factor. Normally when two FortiGate-7000s start up their uptimes are similar and do not affect primary FortiGate-7000 selection. However, during operation, if one of the FortiGate-7000s goes down the other will have a much higher uptime and will be selected as the primary FortiGate-7000 before priority and serial number are tested.

Verifying primary FortiGate-7000 selection

You can use the `diagnose sys ha status` command to verify which FortiGate-7000 has become the primary FortiGate-7000 as shown by the following command output example. This output also shows that the FortiGate-7000 with the highest serial number was selected to be the primary chassis.

```
diagnose sys ha status
==================================================================================
Current slot: 1  Module SN: FIM04E3E16000085
Chassis HA mode: a-p

Chassis HA information:
[Debug_Zone HA information]
HA group member information: is_manage_master=1.
FG74E83E16000015:  Slave, serialno_prio=1, usr_priority=128, hostname=CH15
FG74E83E16000016: Master, serialno_prio=0, usr_priority=127, hostname=CH16
```

How link and module failures affect primary FortiGate-7000 selection

The total number of connected data interfaces in a FortiGate-7000 has a higher priority than the number of failed modules in determining which FortiGate-7000 HA configuration has become the primary FortiGate-7000. For example, if one chassis has a failed FPM and the other has a disconnected or failed data interface, the chassis with the failed FPM becomes the primary unit.

For another example, the following `diagnose sys ha status` command shows the HA status for a cluster where one FortiGate-7000 has a disconnected or failed data interface and the other FortiGate-7000 has a failed FPM.

```
diagnose sys ha status
==========================================================================
Slot: 2 Module SN: FIM01E3E16000088
Chassis HA mode: a-p

Chassis HA information:
[Debug_Zone HA information]
HA group member information: is_manage_master=1.
FG74E33E16000027: Master, serialno_prio=0, usr_priority=128, hostname=Chassis-K
FG74E13E16000072: Slave, serialno_prio=1, usr_priority=128, hostname=Chassis-J

HA member information:
Chassis-K(FIM01E3E16000088), Slave(priority=1), uptime=2237.46, slot=2, chassis=1(1)
    slot: 2, chassis_uptime=2399.58,
    state: worker_failure=1/2, lag=(total/good/down/bad-score)=2/2/0/0,
           intf_state=(port up)=0, force-state(0:none)
           traffic-bandwidth-score=20, mgmt-link=1
    hbdevs: local_interface= 2-M1 best=yes
            local_interface= 2-M2 best=no

Chassis-J(FIM01E3E16000031), Slave(priority=2), uptime=2151.75, slot=2, chassis=2(1)
    slot: 2, chassis_uptime=2151.75,
    state: worker_failure=0/2, lag=(total/good/down/bad-score)=2/2/0/0,
           intf_state=(port up)=0, force-state(0:none)
           traffic-bandwidth-score=20, mgmt-link=1
    hbdevs: local_interface= 2-M1 last_hb_time= 2399.81 status=alive
            local_interface= 2-M2 last_hb_time= 0.00 status=dead

Chassis-J(FIM01E3E16000033), Slave(priority=3), uptime=2229.63, slot=1, chassis=2(1)
    slot: 1, chassis_uptime=2406.78,
    state: worker_failure=0/2, lag=(total/good/down/bad-score)=2/2/0/0,
           intf_state=(port up)=0, force-state(0:none)
           traffic-bandwidth-score=20, mgmt-link=1
    hbdevs: local_interface= 2-M1 last_hb_time= 2399.81 status=alive
            local_interface= 2-M2 last_hb_time= 0.00 status=dead

Chassis-K(FIM01E3E16000086), Master(priority=0), uptime=2203.30, slot=1, chassis=1(1)
    slot: 1, chassis_uptime=2203.30,
    state: worker_failure=1/2, lag=(total/good/down/bad-score)=2/2/0/0,
           intf_state=(port up)=1, force-state(0:none)
           traffic-bandwidth-score=30, mgmt-link=1
    hbdevs: local_interface= 2-M1 last_hb_time= 2399.74 status=alive
            local_interface= 2-M2 last_hb_time= 0.00 status=dead
```

This output shows that chassis 1 (hostname Chassis-K) is the primary or master FortiGate-7000. The reason for this is that chassis 1 has a total `traffic-bandwidth-score` of 30 + 20 = 50, while the total `traffic-bandwidth-score` for chassis 2 (hostname Chassis-J) is 20 + 20 = 40.

The output also shows that both FIMs in chassis 1 have detected a worker failure (`worker_failure=1/2`) while both FIMs in chassis 2 have not detected a worker failure `worker_failure=0/2`). The `intf-state=(port up)=1` field shows that FIM in slot 1 of chassis 1 has one more interface connected than the FIM in slot 1 of chassis 2. It is this extra connected interface that gives the FIM in chassis 1 slot 1 the higher traffic bandwidth score than the FIM in slot 1 of chassis 2.

One of the interfaces on the FIM in slot 1 of chassis 2 must have failed. In a normal HA configuration the FIMs in matching slots of each chassis should have redundant interface connections. So if one module has fewer connected interfaces this indicates a link failure.

FIM failures

If an FIM fails, not only will HA recognize this as a module failure it will also give the chassis with the failed FIM a much lower traffic bandwidth score. So an FIM failure would be more likely to cause an HA failover than a FPM failover.

Also, the traffic bandwidth score for an FIM with more connected interfaces would be higher than the score for an FIM with fewer connected interfaces. So if a different FIM failed in each chassis, the chassis with the functioning FIM with the most connected data interfaces would have the highest traffic bandwidth score and would become the primary chassis.

Management link failures

Management connections to a FortiGate-7000 can affect primary chassis selection. If the management connection to one FortiGate-7000 become disconnected a failover will occur and the FortiGate-7000 that still has management connections will become the primary FortiGate-7000.

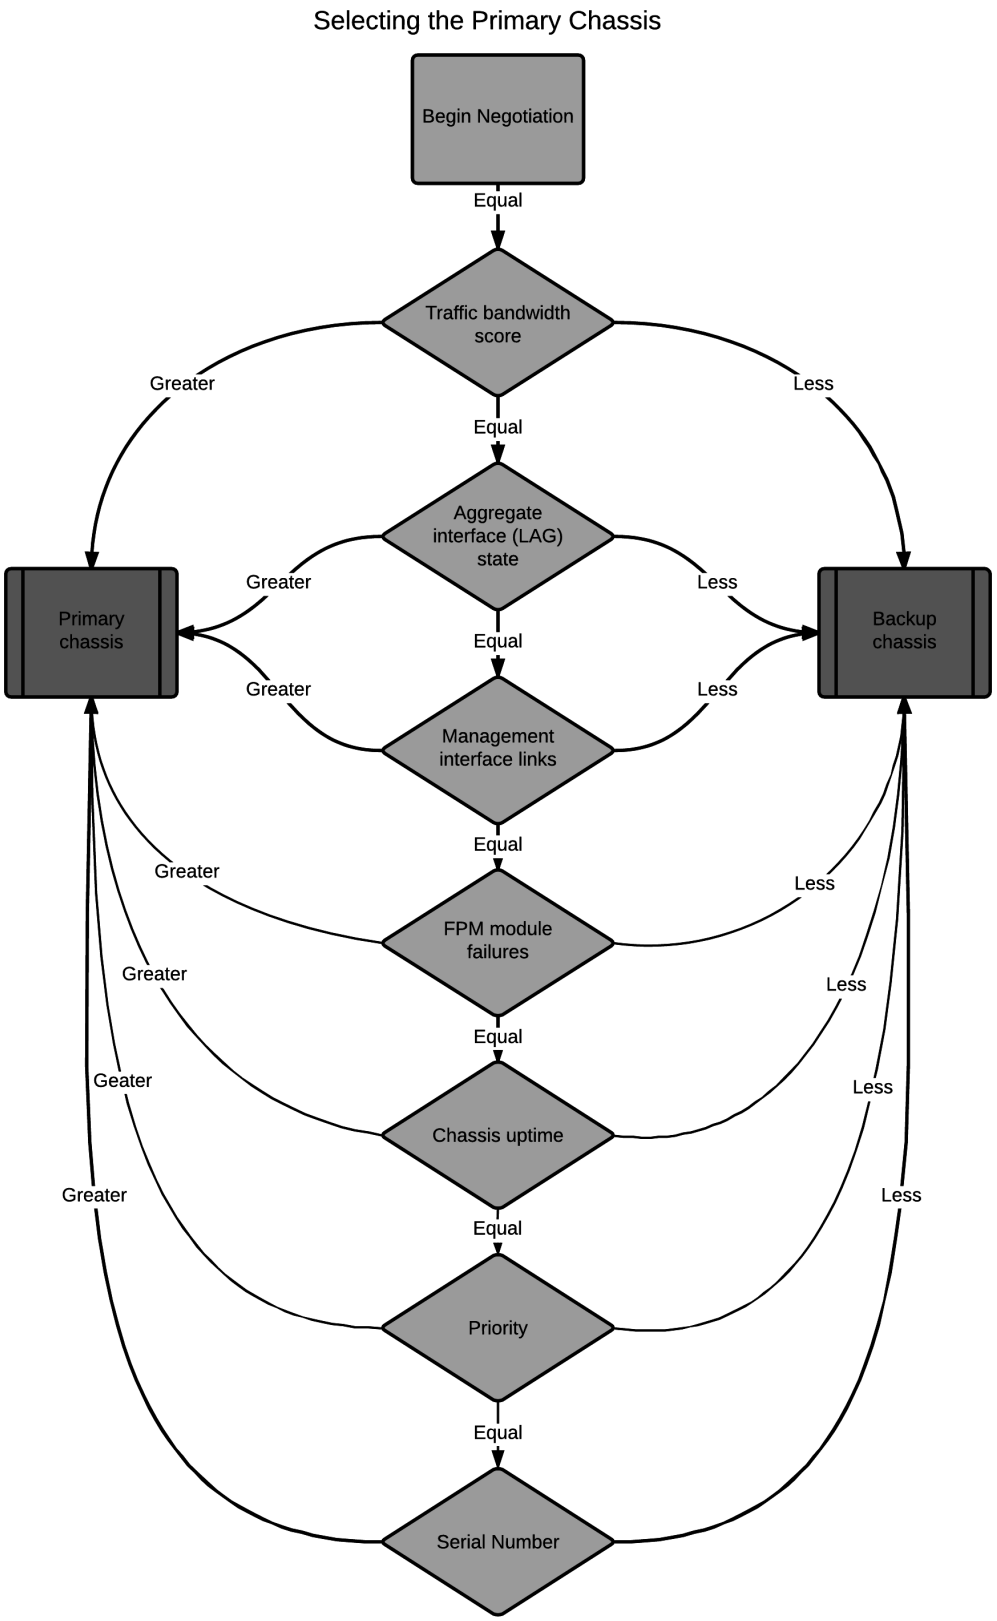
Selecting the Primary Chassis
Begin Negotiation
Equal
Traffic bandwidth score
Greater
Less
Equal
Aggregate interface (LAG) state
Greater
Less
Primary chassis
Backup chassis
Equal
Greater
Less
Management interface links
Equal
Greater
Less
FPM module failures
Greater
Less
Equal
Geater
Less
Chassis uptime
Greater
Less
Equal
Priority
Equal
Serial Number

If there are no failures and if you haven't configured any settings to influence primary chassis selection, the chassis with the highest serial number to becomes the primary chassis.

Using the serial number is a convenient way to differentiate FortiGate-7000 chassis; so basing primary chassis selection on the serial number is predictable and easy to understand and interpret. Also the chassis with the highest serial number would usually be the newest chassis with the most recent hardware version. In many cases you may not need active control over primary chassis selection, so basic primary chassis selection based on serial number is sufficient.

In some situations you may want have control over which chassis becomes the primary chassis. You can control primary chassis selection by setting the priority of one chassis to be higher than the priority of the other. If you change the priority of one of the chassis, during negotiation, the chassis with the highest priority becomes the primary chassis. As shown above, FortiGate-7000 FGCP selects the primary chassis based on priority before serial number. For more information about how to use priorities, see **How link and module failures affect primary FortiGate-7000 selection on page 80**.

Chassis uptime is also a factor. Normally when two chassis start up their uptimes are similar and do not affect primary chassis selection. However, during operation, if one of the chassis goes down the other will have a much higher uptime and will be selected as the primary chassis before priority and serial number are tested.

Verifying primary chassis selection

You can use the `diagnose sys ha status` command to verify which chassis has become the primary chassis as shown by the following command output example. This output also shows that the chassis with the highest serial number was selected to be the primary chassis.

```
diagnose  sys  ha  status
==================================================================
Current slot: 1  Module SN: FIM04E3E16000085
Chassis HA mode: a-p

Chassis HA information:
[Debug_Zone HA information]
HA group member information: is_manage_master=1.
FG74E83E16000015:  Slave, serialno_prio=1, usr_priority=128, hostname=CH15
FG74E83E16000016: Master, serialno_prio=0, usr_priority=127, hostname=CH16
```

Link failure threshold and board failover tolerance

The default settings of the link failure threshold and the board failover tolerance result in the default link and module failure behavior. You can change these settings if you want to modify this behavior. For example, if you want a failover to occur if an FPM fails, even if an interface has failed, you can increase the board failover tolerance setting.

Link failure threshold

The link failure threshold determines how many interfaces in a link aggregation interface (LAG) can be lost before the LAG interface is considered down. The FortiGate-7000 with the most connected LAGs becomes the primary FortiGate-7000. If a LAG goes down the cluster will negotiate and may select a new primary FortiGate-7000. You can use the following command to change the link failure threshold:

```
config system ha
```

```
   set link-failure-threshold <threshold>
end
```

The threshold range is 0 to 80 and 0 is the default.

A threshold of 0 means that if a single interface in any LAG fails the LAG is considered down. A higher failure threshold means that more interfaces in a LAG can fail before the LAG is considered down. For example, if the threshold is set to 1, at least two interfaces will have to fail.

Board failover tolerance

You can use the following command to configure board failover tolerance.

```
config system ha
   set board-failover-tolerance <tolerance>
end
```

The tolerance range is 0 to 12, and 0 is the default.

A tolerance of 0 means that if a single module fails in the primary FortiGate-7000 a failover occurs and the FortiGate-7000 with the fewest failed modules becomes the new primary FortiGate-7000. A higher failover tolerance means that more modules can fail before a failover occurs. For example, if the tolerance is set to 1, at least two modules in the primary FortiGate-7000 will have to fail before a failover occurs.

Priority and primary FortiGate-7000 selection

You can select a FortiGate-7000 to become the primary FortiGate-7000 by setting the HA priority of one or more of its FIMs (for example, the FIM in slot 1) higher than the priority of the other FIMs. Enter the following command to set the HA priority:

```
config system ha
   set priority <number>
end
```

The default priority is 128.

The FortiGate-7000 with the highest total FIM priority becomes the primary FortiGate-7000.

Override and primary FortiGate-7000 selection

Enabling `override` changes the order of primary FortiGate-7000 selection. If you have enabled `override`, primary FortiGate-7000 selection considers priority before up time and serial number. This means that if you set the device priority higher for one FortiGate-7000, with `override` enabled this FortiGate-7000 becomes the primary FortiGate-7000 even if its uptime and serial number are lower than the other FortiGate-7000.

Enter the following command to enable override.

```
config system ha
   set override enable
end
```

When `override` is enabled primary FortiGate-7000 selection checks the traffic bandwidth score, aggregate interface state, management interface links and FPM failures first. Any of these factors can affect primary chassis selection, even if you have enabled override.

FortiGate-7000 v5.4.9 special features and limitations

This section describes special features and limitations for FortiGate-7000 v5.4.9.

Managing the FortiGate-7000

Management is only possible through the MGMT1 to MGMT4 front panel management interfaces. By default the MGMT1 to MGMT4 interfaces of the FIMs in slot 1 and slot 2 are in a single static aggregate interface named mgmt with IP address 192.168.1.99. You manage the FortiGate-7000 by connecting any one of these eight interfaces to your network, opening a web browser and browsing to https://192.168.1.99.

The FortiGate-7030E has one FIM and the MGMT1 to MGMT4 interfaces of that module are the only ones in the aggregate interface.

Default management VDOM

By default the FortiGate-7000 configuration includes a management VDOM named dmgmt-vdom. For the FortiGate-7000 system to operate normally you should not change the configuration of this VDOM and this VDOM should always be the management VDOM. You should also not add or remove interfaces from this VDOM.

You have full control over the configurations of other FortiGate-7000 VDOMs.

Maximum number of LAGs

FortiGate-7000 systems support up to 16 link aggregation groups (LAGs). This includes both normal link aggregation groups and redundant interfaces and including the redundant interface that contains the M1 to M4 management interfaces.

Firewall

TCP sessions with NAT enabled that are expected to be idle for more than the distributed processing normal TCP timer (which is 3605 seconds) should only be distributed to the master FPM using a flow rule. You can configure the distributed normal TCP timer using the following command:

```
config system global
   set dp-tcp-normal-timer <timer>
end
```

UDP sessions with NAT enabled that are expected to be idle for more than the distributed processing normal UDP timer should only be distributed to the primary FPM using a flow rule.

IP multicast

IPv4 and IPv6 Multicast traffic is only sent to the primary FPM (usually the FPM in slot 3). This is controlled by the following configuration:

```
config load-balance flow-rule
   edit 18
      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 19
      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
```

High availability

Only the M1 and M2 interfaces are used for the HA heartbeat communication.

If you use the same switch for both M1 and M2, separate the M1 and M2 traffic on the switch and set the heartbeat traffic on the M1 and M2 Interfaces to have different VLAN IDs. For example, use the following command to set the heartbeat traffic on M1 to use VLAN ID 777 and the heartbeat traffic on M2 to use VLAN ID 888:

```
config system ha
   set hbdev-vlan-id 777
   set hbdev-second-vlan-id 888
end
```

If you don't set different VLAN IDs for the M1 and M2 heartbeat packets, q-in-q must be enabled on the switch.

The following FortiOS HA features are not supported or are supported differently by FortiGate-7000 v5.4.5:

- Remote IP monitoring (configured with the option `pingserver-monitor-interface` and related settings) is not supported.
- Active-active HA is not supported.
- The range for the HA `group-id` is 0 to 14.
- Failover logic for FortiGate-7000 v5.4.5 HA is not the same as FGCP for other FortiGate clusters.
- HA heartbeat configuration is specific to FortiGate-7000 systems and differs from standard HA.
- FortiGate Session Life Support Protocol (FGSP) HA (also called standalone session synchronization) is not supported.
- FortiGate-7000 HA does not support the `route-ttl`, `route-wait`, and `route-hold` options for tuning route synchronization between FortiGate-7000s.

Shelf manager module

It is not possible to access SMM CLI using Telnet or SSH. Only console access is supported using the chassis front panel console ports as described in the FortiGate-7000 system guide.

For monitoring purpose, IPMI over IP is supported on SMM Ethernet ports. See your FortiGate-7000 system guide for details.

FortiOS features not supported by FortiGate-7000 v5.4.9

The following mainstream FortiOS 5.4.5 features are not supported by the FortiGate-7000 v5.4.9:

- Hardware switch
- Usage-based ECMP load balancing is not supported. If the `config system settings v4-ecmp-mode` option is set to `usage-based`, all traffic uses the first ECMP route instead of being load balanced among all ECMP routes. All other ECMP load balancing options are supported, including `source-ip-based`, `weight-based`, and `source-dest-ip-based`.
- Switch controller
- WiFi controller
- WAN load balancing (SD-WAN)
- IPv4 over IPv6, IPv6 over IPv4, IPv6 over IPv6 features
- 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
```

- Hard disk features including, WAN optimization, web caching, explicit proxy content caching, disk logging, and GUI-based packet sniffing.

- The FortiGate-7000 platform only supports quarantining files to FortiAnalyzer.
- Log messages should be sent only using the management aggregate interface

IPsec VPN tunnels terminated by the FortiGate-7000

For a list of new FortiOS 5.4 FortiGate-7000 IPsec VPN features and a list of IPsec VPN features not supported by FortiOS 5.4 FortiGate-7000 IPsec VPN, see FortiGate-7000 IPsec VPN on page 57.

SSL VPN

Sending all SSL VPN sessions to the primary FPM is recommended. You can do this by:

- Creating a flow rule that sends all sessions that use the SSL VPN destination port and IP address to the primary FPM.
- Creating flow rules that send all sessions that use the SSL VPN IP pool addresses to the primary FPM.

Traffic shaping and DDoS policies

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

Sniffer mode (one-arm sniffer)

One-arm sniffer mode is only supported after creating a load balance flow rule to direct sniffer traffic to a specific FPM.

FortiGuard web filtering

All FortiGuard rating queries are sent through management aggregate interface from the management VDOM (named dmgmt-vdom).

Log messages include a slot field

An additional "slot" field has been added to log messages to identify the FPM that generated the log.

FortiOS Carrier

You have to apply a FortiOS Carrier license separately to each FIM and FPM to license a FortiGate-7000 chassis for FortiOS Carrier.

Special notice for new deployment connectivity testing

Only the primary FPM can successfully ping external IP addresses. During a new deployment, while performing connectivity testing from the Fortigate-7000, make sure to run `execute ping` tests from the primary FPM CLI.

Dedicated management interfaces not supported

The FortiGate-7000 does not support configuring dedicated management interfaces using the `config system dedicated-mgmt` command.

FortiGate-7000 v5.4.5 special features and limitations

This section describes special features and limitations for FortiGate-7000 v5.4.5.

Managing the FortiGate-7000

Management is only possible through the MGMT1 to MGMT4 front panel management interfaces. By default the MGMT1 to MGMT4 interfaces of the FIMs in slot 1 and slot 2 are in a single static aggregate interface named mgmt with IP address 192.168.1.99. You manage the FortiGate-7000 by connecting any one of these eight interfaces to your network, opening a web browser and browsing to https://192.168.1.99.

The FortiGate-7030E has one FIM and the MGMT1 to MGMT4 interfaces of that module are the only ones in the aggregate interface.

Default management VDOM

By default the FortiGate-7000 configuration includes a management VDOM named dmgmt-vdom. For the FortiGate-7000 system to operate normally you should not change the configuration of this VDOM and this VDOM should always be the management VDOM. You should also not add or remove interfaces from this VDOM.

You have full control over the configurations of other FortiGate-7000 VDOMs.

Firewall

TCP sessions with NAT enabled that are expected to be idle for more than the distributed processing normal TCP timer (which is 3605 seconds) should only be distributed to the master FPM using a flow rule. You can configure the distributed normal TCP timer using the following command:

```
config system global
   set dp-tcp-normal-timer <timer>
end
```

UDP sessions with NAT enabled that are expected to be idle for more than the distributed processing normal UDP timer should only be distributed to the primary FPM using a flow rule.

IP multicast

IPv4 and IPv6 Multicast traffic is only sent to the primary FPM (usually the FPM in slot 3). This is controlled by the following configuration:

```
config load-balance flow-rule
   edit 18
      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 19
      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
```

High availability

Only the M1 and M2 interfaces are used for the HA heartbeat communication.

If you use the same switch for both M1 and M2, separate the M1 and M2 traffic on the switch and set the heartbeat traffic on the M1 and M2 Interfaces to have different VLAN IDs. For example, use the following command to set the heartbeat traffic on M1 to use VLAN ID 777 and the heartbeat traffic on M2 to use VLAN ID 888:

```
config system ha
   set hbdev-vlan-id 777
   set hbdev-second-vlan-id 888
end
```

If you don't set different VLAN IDs for the M1 and M2 heartbeat packets, q-in-q must be enabled on the switch.

The following FortiOS HA features are not supported or are supported differently by FortiGate-7000 v5.4.5:

- Remote IP monitoring (configured with the option `pingserver-monitor-interface` and related settings) is not supported.
- Active-active HA is not supported.
- The range for the HA `group-id` is 0 to 14.
- Failover logic for FortiGate-7000 v5.4.5 HA is not the same as FGCP for other FortiGate clusters.

- HA heartbeat configuration is specific to FortiGate-7000 systems and differs from standard HA.
- FortiGate Session Life Support Protocol (FGSP) HA (also called standalone session synchronization) is not supported.
- FortiGate-7000 HA does not support the `route-ttl`, `route-wait`, and `route-hold` options for tuning route synchronization between FortiGate-7000s.

Shelf manager module

It is not possible to access SMM CLI using Telnet or SSH. Only console access is supported using the chassis front panel console ports as described in the FortiGate-7000 system guide.

For monitoring purpose, IPMI over IP is supported on SMM Ethernet ports. See your FortiGate-7000 system guide for details.

FortiOS features not supported by FortiGate-7000 v5.4.5

The following mainstream FortiOS 5.4.5 features are not supported by the FortiGate-7000 v5.4.5:

- Hardware switch
- Usage-based ECMP load balancing is not supported. If the `config system settings v4-ecmp-mode` option is set to `usage-based`, all traffic uses the first ECMP route instead of being load balanced among all ECMP routes. All other ECMP load balancing options are supported, including `source-ip-based`, `weight-based`, and `source-dest-ip-based`.
- Switch controller
- WiFi controller
- WAN load balancing (SD-WAN)
- IPv4 over IPv6, IPv6 over IPv4, IPv6 over IPv6 features
- 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
```

- Hard disk features including, WAN optimization, web caching, explicit proxy content caching, disk logging, and GUI-based packet sniffing.
- The FortiGate-7000 platform only supports quarantining files to FortiAnalyzer.
- Log messages should be sent only using the management aggregate interface

IPsec VPN tunnels terminated by the FortiGate-7000

This section lists FortiGate-7000 limitations for IPsec VPN tunnels terminated by the FortiGate-7000:

- Interface-based IPsec VPN is recommended.
- Policy based IPsec VPN is supported, but requires creating flow-rules for each Phase 2 selector.
- Dynamic routing and policy routing is not supported for IPsec interfaces.
- IPsec static routes don't consider distance, weight, priority settings. IPsec static routes are always installed in the routing table, regardless of the tunnel state.
- IPsec tunnels are not load-balanced across the FPMs, all IPsec tunnel sessions are sent to the primary FPM.
- IPsec VPN dialup or dynamic tunnels require a flow rule that sends traffic destined for IPsec dialup IP pools to the primary FPM.
- In an HA configuration, IPsec SAs are not synchronized to the backup chassis. IPsec SAs are renegotiated after a failover.
- Auto-negotiate is not supported for IPsec VPN phase 2 configurations. IPsec VPN tunnels should not be started by the FortiGate-7000 because this could lead to every FPM attempting to start an IPsec VPN tunnel with the same remote device.

SSL VPN

Sending all SSL VPN sessions to the primary FPM is recommended. You can do this by:

- Creating a flow rule that sends all sessions that use the SSL VPN destination port and IP address to the primary FPM.
- Creating flow rules that send all sessions that use the SSL VPN IP pool addresses to the primary FPM.

Traffic shaping and DDoS policies

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

Sniffer mode (one-arm sniffer)

One-arm sniffer mode is only supported after creating a load balance flow rule to direct sniffer traffic to a specific FPM.

FortiGuard web filtering

All FortiGuard rating queries are sent through management aggregate interface from the management VDOM (named dmgmt-vdom).

Log messages include a slot field

An additional "slot" field has been added to log messages to identify the FPM that generated the log.

FortiOS Carrier

You have to apply a FortiOS Carrier license separately to each FIM and FPM to license a FortiGate-7000 chassis for FortiOS Carrier.

Special notice for new deployment connectivity testing

Only the primary FPM can successfully ping external IP addresses. During a new deployment, while performing connectivity testing from the Fortigate-7000, make sure to run `execute ping` tests from the primary FPM CLI.

Dedicated management interfaces not supported

The FortiGate-7000 does not support configuring dedicated management interfaces using the `config system dedicated-mgmt` command.

FortiGate-7000 v5.4.3 special features and limitations

This section describes special features and limitations for FortiGate-7000 v5.4.3.

Managing the FortiGate-7000

Management is only possible through the MGMT1 to MGMT4 front panel management interfaces. By default the MGMT1 to MGMT4 interfaces of the FIMs in slot 1 and slot 2 are in a single static aggregate interface named mgmt with IP address 192.168.1.99. You manage the FortiGate-7000 by connecting any one of these eight interfaces to your network, opening a web browser and browsing to https://192.168.1.99.

The FortiGate-7030E has one FIM and the MGMT1 to MGMT4 interfaces of that module are the only ones in the aggregate interface.

Default management VDOM

By default the FortiGate-7000 configuration includes a management VDOM named dmgmt-vdom. For the FortiGate-7000 system to operate normally you should not change the configuration of this VDOM and this VDOM should always be the management VDOM. You should also not add or remove interfaces from this VDOM.

You have full control over the configurations of other FortiGate-7000 VDOMs.

Firewall

TCP sessions with NAT enabled that are expected to be idle for more than the distributed processing normal TCP timer (which is 3605 seconds) should only be distributed to the master FPM using a flow rule. You can configure the distributed normal TCP timer using the following command:

```
config system global
   set dp-tcp-normal-timer <timer>
end
```

UDP sessions with NAT enabled that are expected to be idle for more than the distributed processing normal UDP timer should only be distributed to the primary FPM using a flow rule.

Link monitoring and health checking

ICMP-based link monitoring for SD-WAN, ECMP, HA link monitoring, and firewall session load balancing monitoring (or health checking) is not supported. Using TCP or UDP options for link monitoring instead.

IP multicast

IPv4 and IPv6 Multicast traffic is only sent to the primary FPM (usually the FPM in slot 3). This is controlled by the following configuration:

```
config load-balance flow-rule
   edit 18
      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 19
      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
```

High availability

Only the M1 and M2 interfaces are used for the HA heartbeat communication.

When using both M1 and M2 for the heartbeat, FortiGate-7000 v5.4.3 requires two switches. The first switch to connect all M1 ports together. The second switch to connect all M2 ports together. This is because the same VLAN is used for both M1 and M2 and the interface groups should remain in different broadcast domains.

Using a single switch for both M1 and M2 heartbeat traffic is possible if the switch supports q-in-q tunneling. In this case use different VLANs for M1 traffic and M2 traffic to keep two separated broadcast domains in the switch.

The following FortiOS HA features are not supported or are supported differently by FortiGate-7000 v5.4.3:

- Remote IP monitoring (configured with the option `pingserver-monitor-interface` and related settings) is not supported
- Active-active HA is not supported
- The range for the HA `group-id` is 0 to 14.
- Failover logic for FortiGate-7000 v5.4.3 HA is not the same as FGCP for other FortiGate clusters.
- HA heartbeat configuration is specific to FortiGate-7000 systems and differs from standard HA.
- FortiGate Session Life Support Protocol (FGSP) HA (also called standalone session synchronization) is not supported.

Shelf manager module

It is not possible to access SMM CLI using Telnet or SSH. Only console access is supported using the chassis front panel console ports as described in the FortiGate-7000 system guide.

For monitoring purpose, IPMI over IP is supported on SMM Ethernet ports. See your FortiGate-7000 system guide for details.

FortiOS features not supported by FortiGate-7000 v5.4.3

The following mainstream FortiOS 5.4.3 features are not supported by the FortiGate-7000 v5.4.3:

- Hardware switch
- Usage-based ECMP load balancing is not supported. If the `config system settings v4-ecmp-mode` option is set to `usage-based`, all traffic uses the first ECMP route instead of being load balanced among all ECMP routes. All other ECMP load balancing options are supported, including `source-ip-based`, `weight-based`, and `source-dest-ip-based`.
- Switch controller
- WiFi controller
- WAN load balancing (SD-WAN)
- IPv4 over IPv6, IPv6 over IPv4, IPv6 over IPv6 features
- 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
```

- Hard disk features including, WAN optimization, web caching, explicit proxy content caching, disk logging, and GUI-based packet sniffing.
- The FortiGate-7000 platform only supports quarantining files to FortiAnalyzer.
- Log messages should be sent only using the management aggregate interface

IPsec VPN tunnels terminated by the FortiGate-7000

This section lists FortiGate-7000 limitations for IPsec VPN tunnels terminated by the FortiGate-7000:

- Interface-based IPsec VPN is recommended.
- Policy based IPsec VPN is supported, but requires creating flow-rules for each Phase 2 selector.
- Dynamic routing and policy routing is not supported for IPsec interfaces.
- Remote network subnets are limited to /16 prefix.
- IPsec static routes don't consider distance, weight, priority settings. IPsec static routes are always installed in the routing table, regardless of the tunnel state.
- IPsec tunnels are not load-balanced across the FPMs, all IPsec tunnel sessions are sent to the primary FPM.
- IPsec VPN dialup or dynamic tunnels require a flow rule that sends traffic destined for IPsec dialup IP pools to the primary FPM.
- In an HA configuration, IPsec SAs are not synchronized to the backup chassis. IPsec SAs are renegotiated after a failover.
- Auto-negotiate is not supported for IPsec VPN phase 2 configurations. IPsec VPN tunnels should not be started by the FortiGate-7000 because this could lead to every FPM attempting to start an IPsec VPN tunnel with the same remote device.

More about IPsec VPN routing limitations

For IPv4 traffic, FortiGate-7000s can only recognize netmasks with 16-bit or 32-bit netmasks. For example:

The following netmasks are supported:

- 12.34.0.0/24
- 12.34.0.0 255.255.0.0
- 12.34.56.0/21
- 12.34.56.0 255.255.248.0
- 12.34.56.78/32
- 12.34.56.78 255.255.255.255
- 12.34.56.78 (for single IP addresses, FortiOS automatically uses 32-bit netmasks)

The following netmasks are not supported:

- 12.34.0.0/15 (netmask is less than 16-bit)
- 12.34.0.0 255.254.0.0 (netmask is less than 16-bit)
- 12.34.56.1-12.34.56.100 (ip range is not supported)
- 12.34.56.78 255.255.220.0 (invalid netmask)

SSL VPN

Sending all SSL VPN sessions to the primary FPM is recommended. You can do this by:

- Creating a flow rule that sends all sessions that use the SSL VPN destination port and IP address to the primary FPM.
- Creating flow rules that send all sessions that use the SSL VPN IP pool addresses to the primary FPM.

Authentication

This section lists FortiGate-7000 authentication limitations:

- Active authentication that requires a user to manually log into the FortiGate firewall can be problematic because the user may be prompted for credentials more than once as sessions are distributed to different FPMs. You can avoid this by changing the load distribution method to `src-ip`.
- FSSO is supported. Each FPM independently queries the server for user credentials.
- RSSO is only supported after creating a load balance flow rule to broadcast RADIUS accounting messages to all FPMs.

Traffic shaping and DDoS policies

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

Sniffer mode (one-arm sniffer)

One-arm sniffer mode is only supported after creating a load balance flow rule to direct sniffer traffic to a specific FPM.

FortiGuard Web Filtering

All FortiGuard rating queries are sent through management aggregate interface from the management VDOM (named dmgmt-vdom).

Log messages include a slot field

An additional "slot" field has been added to log messages to identify the FPM that generated the log.

FortiOS Carrier

FortiOS Carrier is supported by the FortiGate-7000 v5.4.3 but GTP load balancing is not supported.

You have to apply a FortiOS Carrier license separately to each FIM and FPM to license a FortiGate-7000 chassis for FortiOS Carrier.

Special notice for new deployment connectivity testing

Only the primary FPM can successfully ping external IP addresses. During a new deployment, while performing connectivity testing from the Fortigate-7000, make sure to run `execute ping` tests from the primary FPM CLI.

Dedicated management interfaces not supported

The FortiGate-7000 does not support configuring dedicated management interfaces using the `config system dedicated-mgmt` command.

FortiGate-7000 load balancing commands

The most notable difference between a FortiGate-7000 and other FortiGates are the commands described in this section for configuring load balancing. The following commands are available:

```
config load-balance flow-rule
config load-balance setting
```

In most cases you do not have to use these commands. However, they are available to customize some aspects of load balancing.

FortiGate-6000 and FortiGate-7000 have the commands described in this chapter in common. Some of the options described in this section may not apply in the same way on these two platforms. These differences are noted in the descriptions. For example, the generic term "worker" maps to an FPC in a FortiGate-6000 and an FPM in a FortiGate-7000.

config load-balance flow-rule

Use this command to add FortiGate-6000 or FortiGate-7000 flow rules that add exceptions to how matched traffic is processed. Specifically you can use these 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 0
      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 {any | icmp | icmpv6 | tcp | udp | igmp | sctp | gre | esp | ah | ospf | pim
           | vrrp}
      set src-l4port <start>[-<end>]
      set dst-l4port <start>[-<end>]
      set tcp-flag
      set action {forward | mirror-ingress stats | drop}
      set mirror-interface <interface-name>
      set forward-slot {master | all | load-balance | <FPM# or FPC#>}
      set priority <number>
      set comment <text>
   end
```

status {disable | enable}

Enable or disable this flow rule. Default for a new flow-rule is disable.

src-interface <interface-name> [interface-name>...}

The names of one or more FIM interface front panel interfaces accepting the traffic to be subject to the flow rule.

vlan <vlan-id>

If the traffic matching the rule is VLAN traffic, enter the VLAN ID used by the traffic.

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, or IPv4 traffic.

{src-addr-ipv4 | dst-addr-ipv4} <ipv4-address> <netmask>

The source and destination address of the traffic to be matched. The default of `0.0.0.0 0.0.0.0` matches all traffic. Available if `ether-type` is set to `ipv4`.

{src-addr-ipv6 | dst-addr-ipv6} <ip-address> <netmask>

The source and destination address of the traffic to be matched. The default of `::/0` matches all traffic. Available if `ether-type` is set to `ipv6`.

protocol {any | icmp | 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`.

{src-l4port | dst-l4port} <start>[-<end>]

Specify a source port range and a destination port range. This option appears for some protocol settings. For example if protocol is set to tcp or udp. The default range is 0-0.

action {forward | mirror-ingress | stats | drop}

How to handle matching packets. 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 add multiple options.

The default action is `forward`.

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 when `action` is set to `mirror-ingress` or `mirror-egress`.

forward-slot {master | all | load-balance | <FPC#>}

The slot that you want to forward the traffic that matches this rule to.

`master` forwards the primary FPM or FPC.

`all` means forward the traffic to all slots.

`load-balance` means use the default load balancing configuration to handle this traffic.

`<FPM# or FPC#>` allows you to forward the matching traffic to a specific FPC (FortiGate-6000) or FPM module (FortiGate-7000). For example, FPC3 is the FortiGate-6000 FPC in slot 3. FPM3 is the FortiGate-7000 FPM module in slot 3.

priority <number>

Set the priority of the flow rule in the range 1 (highest priority) to 10 (lowest priority). Higher priority rules are matched first. You can use the priority to control which rule is matched first if you have overlapping rules.

comment <text>

Optionally add a comment that describes the 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 ipsec-load-balance {disable | enable}
   set gtp-load-balance {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}
   config workers
      edit 3
         set status enable
         set weight 5
      end
end
```

slbc-mgmt-intf {mgmt1 | mgmt2 | mgmt3}

Select the interface used for management connections. The name depends on the available management interfaces on the device you are managing. You can use this option to enable access to the FPCs directly using the special administration ports as described in Managing individual FPCs on page 1.

max-miss-heartbeats <heartbeats>

Set the number of missed heartbeats before a worker (an FPC or FPM module) is considered to have failed. If this many heartbeats are not received from a worker, this indicates that the worker is not able to process data traffic and no more traffic will be sent to this worker.

The time between heartbeats is 0.2 seconds. Range is 3 to 300. 3 means 0.6 seconds, 10 (the default) means 2 seconds, and 300 means 60 seconds.

max-miss-mgmt-heartbeats <heartbeats>

Set the number of missed management heartbeats before a worker (an FPC or FPM module) is considering to have failed. If this happens, there is a communication problem between a worker and other workers. This communication problem means the worker may not be able to synchronize configuration changes, sessions, the kernel routing table, the bridge table and so on with other workers. If a management heartbeat failure occurs, no traffic will be sent to the worker.

The time between management heartbeats is 1 second. Range is 3 to 300 seconds. The default is 20 seconds.

weighted-load-balance {disable | enable}

Enable weighted load balancing depending on the slot (or worker) weight. Use the `config workers` command to set the weight for each slot or worker.

gtp-load-balance {disable | enable}

Enable GTP load balancing. If GTP load balancing is enabled, Tunnel Endpoint Identifiers (TEIDs) are used to identify GTP sessions.

dp-load-distribution-method {to-master | 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 distribute sessions among workers (FPCs or FPM modules). Usually you would only need to change the 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-ipsport-dport` sessions are distributed 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. All session information is taken 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.

config workers

Set the weight and enable or disable each worker (FPC or FPM module). Use the edit command to specify the slot the worker is installed in. You can enable or disable each worker and set each worker's weight.

The weight range is 1 to 10. 5 is average, 1 is -80% of average and 10 is +100% of average. The weights take effect if `weighted-loadbalance` is enabled.

```
config workers
   edit 3
      set status enable
      set weight 5
   end
```