



FortiADC - HA Deployment Guide

Version 7.2.0

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/training-certification>

FORTINET TRAINING INSTITUTE

<https://training.fortinet.com>

FORTIGUARD CENTER

<https://www.fortiguard.com>

END USER LICENSE AGREEMENT

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

FEEDBACK

Email: techdoc@fortinet.com



February 3, 2023

FortiADC 7.2.0 HA Deployment Guide

01-544-677187-20230203

TABLE OF CONTENTS

Change Log	4
About this guide	5
HA modes	6
HA Active-Passive Mode	6
HA Active-Active Mode	6
HA VRRP Mode	7
Choose HA mode	8
HA deployment	9
Deploy HA-AP mode	9
Deploy HA-AA mode	14
Deploy HA-VRRP mode	21
HA troubleshooting	30
HA management interface	30
HA config out of sync	30
HA on Microsoft Hyper-V platform	31
HA abnormal state	32
Upgrade Firmware	32
HA debug	33

Change Log

Date	Change Description
2020-04-21	Initial release.

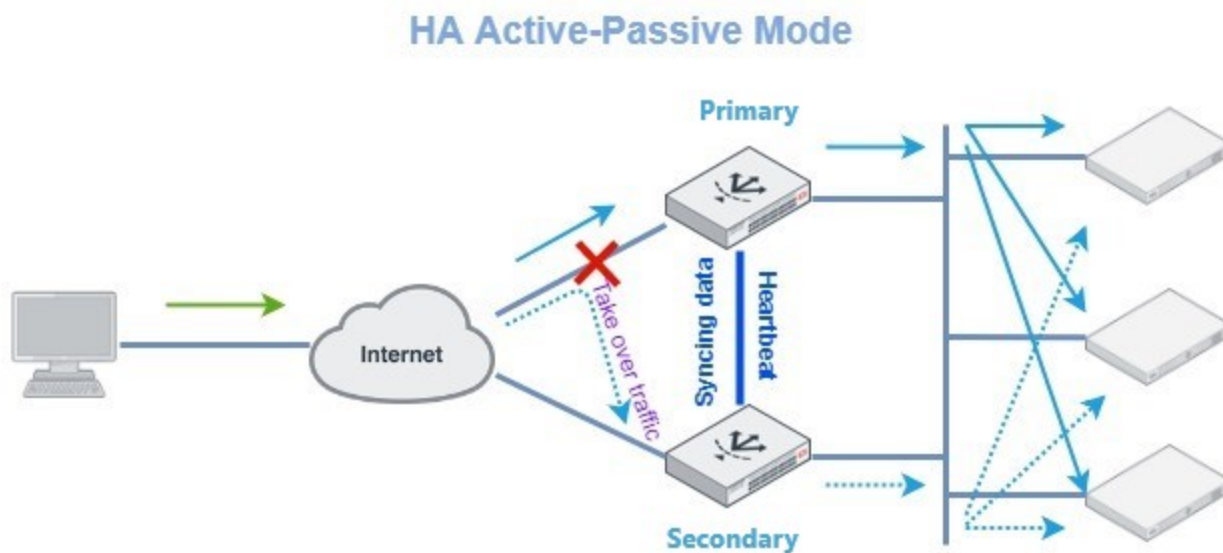
About this guide

This guide details the steps required to configure the FortiADC HA (High Availability) mode. HA aims to ensure an agreed level of operational performance, usually uptime, for a higher than normal period. FortiADC device can be deployed as single units or as a clustered pair. We always recommend deploying a clustered pair to avoid introducing a single point of failure.

HA modes

HA Active-Passive Mode

When the FortiADC devices are configured as HA Active-Passive mode, the active device (also called Primary) handles all the traffic under normal circumstances. If something wrong happens on the active device, the passive device (also called Secondary) becomes active and handles all the traffic instead.



Above chart is the HA-AP mode deployment. Normally, Secondary doesn't handle the traffic, all the traffic is handled by the Primary whatever for the client side or server side. However, the Secondary can always sync the data from Primary, such as incremental configuration changes, layer4 session/persistence table, layer7 persistence, health-check status. Once there is something wrong with the current Primary, such as the monitored interfaces are down (in this case the monitored interface connected to ISP directly), or the physical device is failing, the Secondary will become the new Primary, so that handle all the traffic going through FortiADC.

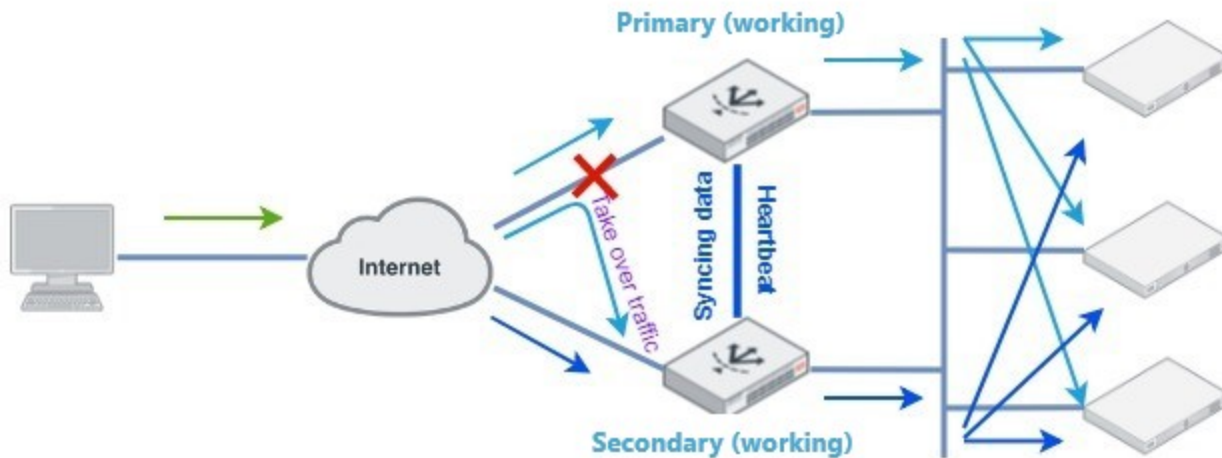
HA-AP mode is the most stable deployment mode, and it can be deployed on any platforms without problem. In this mode, the FortiADC's interface is applied with virtual mac address, once the HA peer takes over the Primary, new Primary will inherit the virtual mac address on the interfaces. This can reduce the traffic failing time while failover happening. On the other hand, this can provide the benefit if the security device such as firewall in your network need Mac address binding. Please be aware that HA-AP mode on Microsoft Hyper-V platform uses the physical Mac Address due to the platform limitation.

HA Active-Active Mode

In the HA Active-Active mode, both the Primary and Secondary is able to handle the traffic normally. There is one thing should be detailed. Although both Primary and Secondary can handle the traffic, FortiADC can only sync the layer4 virtual-server session to its peers. So for layer4 traffic, if the traffic returned from real server goes to FortiADC devices

which is different from the inbound traffic, this FortiADC can still forward the traffic back to client. For the traffic which will be routed by FortiADC, it has the similar issue. This may cause the performance decrease. So ideally, you should have a routing device between FortiADC and real servers, which has the function can send the return traffic to its original FortiADC devices. For layer 7 virtual-server, usually FortiADC establishes the session to real servers by its own interface IP address, so the traffic can be returned to itself natively, unless you enable the “source-address” on it.

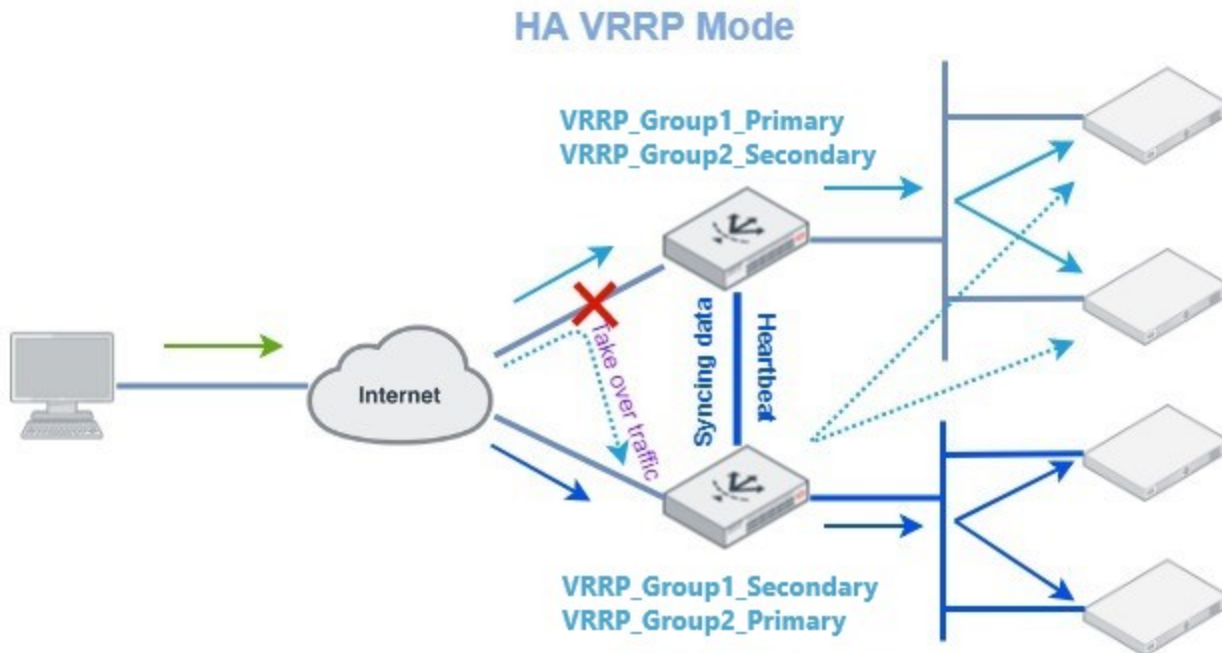
HA Active-Active Mode



If one of the device monitored link is down or if even the entire device is failing, its HA peer can take over all the traffic.

HA VRRP Mode

HA-VRRP mode on the other hand divides the resources into groups, so that we can create multiple VRRP groups, and then assign the public IP resources into the groups. In this way, we can get the another active-active mode. In this mode, every HA node has its own interface IP, but we can define the floating IP on the interface which belongs to one of the VRRP groups.



In general, the connected devices or servers are pointing the gateway to the floating-ip of the VRRP group. The floating-ip is the IP address which can only work on the VRRP group Primary. If the failover happens, the floating-ip will work on the new VRRP Primary. This can assure the floating-ip is always online.

Look at the chart above, this is an example of HA-VRRP mode. Typically, we create 2 VRRP groups, let's say, one is the VRRP_Group1, the other one is VRRP_Group2. We make FortiADC1 the Primary of VRRP_Group1, the Secondary of VRRP_Group2; while the FortiADC2 will be the Secondary of VRRP_Group1, and Primary of VRRP_Group2. Then we divide the real servers into these 2 groups. The servers in group1 point the default gateway to VRRP_Group1's floating-ip, while the servers in group2 point the default gateway to VRRP_Group2's floating-ip. Then normally, FortiADC1 handles the traffic to VRRP_Group1, FortiADC2 handles the traffic to VRRP_Group2. If one of the monitored link or device is down, the HA peer can take over the traffic.

Choose HA mode

We support 3 kinds of HA deployment mode. They are HA-AP, HA-AA, HA-VRRP mode. If you are willing to have a very stable system, please use the HA-AP mode. Although only one device is processing the data, the backup device can take over the Primary work smoothly, it offers the most reliable environment while needs least deployment conditions.

If you are interested in the all active plan, then the HA-VRRP should be your first choice. Once you've arranged the servers in group, all the FortiADC devices can handle the traffic, which increase the throughput and other performance a lot. It requires less deployment conditions over HA-AA mode, and can provide the performance increasing.

Only if you can make sure you can make all the ideal preconditions for the HA-AA, then choose the HA-AA mode. This mode requires the most, but providing the easier configuration logic. It can also offer the good performance in the ideal conditions.

HA deployment

This section includes the steps for deploying HA in the following modes:

[Deploy HA-AP mode on page 9](#)

[Deploy HA-AA mode on page 14](#)

[Deploy HA-VRRP mode on page 21](#)

Deploy HA-AP mode

1) Enable the management interface

It is recommended that the management-interface should be enabled when the HA-AP mode is deployed. Because once you complete the HA-AP mode, only master can handle the traffic, it means that you're not able to access Secondary device directly. It is not convenient in most cases. Management-interface on the other hand, is a virtual-interface binding to the physical interface. It can always work on all the modes including standalone. Please perform the following steps on all the HA nodes.

Steps

1. Get the console control for FortiADC, execute the next steps in the console.
2. Since the manage-interface is a virtual-interface inside the system, so it has the similar routing mechanism as other interface. So there should be no overlapping subnet in the system. Therefore, usually we clear the original IP address of the physical interface.

This can result in the losing the connectivity, so the first step is requiring the console.

```
FAD2 # config system interface
FAD2 (interface) # edit port1
FAD2 (port1) # unset ip
FAD2 (port1) # end
```

3. Configure the management interface.

```
FAD2 # config system ha
FAD2 (ha) # set mgmt-status enable
FAD2 (ha) # set mgmt-interface port1
FAD2 (ha) # set mgmt-ip 10.106.188.42/23
FAD2 (ha) # set mgmt-ip-allowaccess http https ping snmp ssh telnet
FAD2 (ha) # end
```

4. Configure the default route accordingly.

```
FAD2 # config router static
FAD2 (static) # edit 1
FAD2 (1) # set gateway 10.106.189.254
FAD2 (1) # end
```

Note:

On the virtualization platform such as VMware ESXi, KVM, Hyper-V and so on. The VM interface which you are going to bind the management-interface should enable the Promiscuous mode. This mode has different name on different platform, for example it is called "MAC address spoofing" on Hyper-V platform.

2) Configure the HA-AP mode on both sides

Once you completed the management-interface, then you can perform the following steps on Web-UI.

Steps

1. Plan the HA role for the devices

There are two types of HA roles you have to plan, one is the traffic-role, the other is the config-role. Technically, you can configure the traffic-Primary and config-Primary on different devices. Only the traffic-Primary can handle the traffic, and the full configuration sync can be only from the config-Primary to others. (Incremental configuration sync can happen from any side).

Typically, the traffic-Primary and config-Primary are the same one. Here is the example to configure the traffic-Primary and config-Primary on the same device with override enabled.

The condition to make sure negotiation successfully:

- All the HA devices use the same heartbeat ports and data ports.
- All the HA devices have same group-id

How the traffic-Primary is elected in HA-AP mode:

Override enabled:

Disk state > monitor interface > priority > uptime > SN

Override disabled:

Disk state > monitor interface > uptime > priority > SN

- Disk state means the harddisk working state, the device without harddisk error wins. If all the devices have disk error, then compare the next condition.
- Monitor interface means the up monitored interfaces count, devices with more up interfaces wins, if all the devices have the same number of up interfaces, then compare the next condition.
- Priority is the value specified in HA configuration, device with lower value wins, if all the devices have same value, then compare next condition.
- Uptime is the uptime of the device, device with long uptime wins, if all the devices have the same uptime, then compare the next condition.
- SN means the serial number, the device with higher SN will be the Primary.

How the config-Primary is elected (This is same in 3 modes):

config-priority > SN

- Config-priority is the value specified in HA config, the device with lower config-priority value will be the config-Primary.
- SN means the serial number, the device with higher SN will be the config-Primary.

Here we set up 2 HA devices running HA-AP mode, make FAD1 the Primary, and the FAD2 the Secondary. We put config example like following.

FAD1:

```
config system ha
  set mode active-passive
  set hbdev port6 port7
  set group-id 14
  set group-name group1
  set priority 1
  set config-priority 10
  set override enable
```

```

set l7-persistence-pickup enable
set l4-persistence-pickup enable
set l4-session-pickup enable
set monitor port2 port3 port4 port5
end

```

FAD2:

```

config system ha
set mode active-passive
set hbdev port6 port7
set group-id 14
set group-name group1
set priority 9
set config-priority 100
set override enable
set l7-persistence-pickup enable
set l4-persistence-pickup enable
set l4-session-pickup enable
set monitor port2 port3 port4 port5
end

```

There are some preconditions for the HA negotiation:

- The hostname of HA nodes must NOT be same
- The group-id of HA nodes must be same
- The heartbeat interfaces should be connected directly or in the same VLAN
- On some virtualization platforms like Hyper-V, the heartbeat interface should enable the “Mac address spoofing”.
- Configure the basic HA options

The following example shows the FAD1 configuration, the FAD2 is similar.

Navigate to “System->High Availability” page:

The screenshot shows the FortiADC FAD1 High Availability configuration page. The left sidebar has a search bar and a menu with options: Dashboard, FortiView, System, Settings, Traffic Group, Administrator, SNMP, Debug, Certificate, Manage Certificates, Verify, Alert, and Alert. The 'Settings' option is highlighted with a red box. The main content area shows the 'High Availability' tab. It displays the 'HA Cluster Status' as 'standalone'. Below this, there is a status bar with indicators for UP, DOWN, Low, High, and Busy. A table lists the host FAD1 with its state, serial number, node ID, IP address, and config source. A red box highlights the edit icon in the table, with a red arrow pointing to it and the text 'Click here to edit the HA'.

Host Name	State	Serial Number	Node ID	IP Address	Config Source
FAD1	Standalone	FADV040000146261	0	169.254.16.97	N/A

Configure the required options.

High Availability Setting

Basic

Synchronization

Advanced

Cluster Mode

Standalone

Active-Passive

Active-Active

Active-Active-VRRP

Group Name

Group1

Group ID

14

Default: 0 Range: 0-31

Config Priority

10

Default: 100 Range: 0-255

Monitor Interface

Heartbeat Interface

* port6

* port7

Data Interface

Save

Cancel

Configure the synchronization options.

High Availability Setting

Basic

Synchronization

Advanced

Layer 7 Persistence Synchronization

ON

Layer 4 Persistence Synchronization

ON

Layer 4 Connection Synchronization

ON

Save

Cancel

Configure the advanced options.

High Availability Setting

Basic Synchronization **Advanced**

Priority
1
Default: 5 Range: 0-9

Override
ON

Heartbeat Interval
2
Default: 2 Range: 1-20 intervals (100 milliseconds per interval)

Lost Heartbeat Threshold
6
Default: 6 Range: 1-60 retries

ARP Times
5
Default: 5 Range: 1-60 times

ARP Interval
6
Default: 6 Range: 1-20 seconds

Save Cancel

Deploy HA-AA mode

1) Plan the HA deployment

The condition to make sure negotiation successfully:

- All the HA devices use the same heartbeat ports and data ports.
- All the HA devices have same group-id
- All the HA devices have same node-list
- All the HA devices have different local-node-id

How the traffic-Primary is elected in HA-AA mode:

Override enabled:

Disk state > monitor interface > Remote IP check > priority > uptime > SN

Override disabled:

Disk state > monitor interface > Remote IP check > uptime > priority > SN

- Disk state means the harddisk working state, the device without harddisk error wins. If all the devices have disk error, then compare the next condition.

- Monitor interface means the up monitored interfaces count, devices with more up interfaces wins, if all the devices have the same number of up interfaces, then compare the next condition.
- Priority is the value specified in HA configuration, device with lower value wins, if all the devices have same value, then compare next condition.
- Uptime is the uptime of the device, device with long uptime wins, if all the devices have the same uptime, then compare the next condition.
- SN means the serial number, the device with higher SN will be the Primary.

How the config-Primary is elected (This is same in 3 modes):

- Config-priority is the value specified in HA config, the device with lower config-priority value will be the config-Primary.
- SN means the serial number, the device with higher SN will be the config-Primary.

Some important notes:

- In HA-AA mode, every device interface has its own working IP address, it should be specified under “config ha-node-ip-list”. These config can be synced to all the HA peers. Each HA peer uses its own IP according to local-node-id.
- Each HA node should have its own local-node-id, and the local-node-id on different nodes must be different.
- In HA-AA mode, the IP address of interface is not working any longer, only the IP address under “config ha-node-ip-list” can work accordingly.
- To achieve the best performance and stable environment, you need to set up a routing device (typically router) between FortiADC and real-servers. The routing device should have the function like “reverse-route”, it means that the return packets from real-servers can be forwarded back to the original FortiADC node which distributed the traffic to the real-server. For example, if the requests from client1 were handled by FortiADC1, the FortiADC1 distributes the requests to real-server1, the return packets from real-server1 to client1 should be forwarded to FortiADC1 back by the routing device.

In this example, we’re going to make FAD1 the traffic-Primary and config-Primary, FAD2 the traffic-Secondary and config-Secondary. If you have management-interface, then you can configure it in Web-UI, otherwise, you’d better configure it from console.

FAD1:

```
config system ha
  set mode active-active
  set hbdev port6 port7
  set group-id 14
  set node-list 0 1
  set group-name group1
  set priority 3
  set config-priority 40
  set override enable
  set l7-persistence-pickup enable
  set l4-persistence-pickup enable
  set l4-session-pickup enable
  set monitor port2 port3 port4 port5
end
```

FAD2:

```
config system ha
  set mode active-active
  set hbdev port6 port7
  set group-id 14
  set node-list 0 1
  set local-node-id 1
```

```

set group-name group1
set priority 9
set config-priority 100
set override enable
set l7-persistence-pickup enable
set l4-persistence-pickup enable
set l4-session-pickup enable
set monitor port2 port3 port4 port5
end

```

2) Configure the HA options

The following example shows the FAD1 configuration, the FAD2 is similar.

Navigate to “System->High Availability” page:

The screenshot shows the FortiADC FAD1 configuration page for High Availability. The left sidebar has a red box around 'High Availability'. The main content area shows 'HA Cluster Status' for FAD1 in 'Standalone' mode. Below this is a table with columns: Host Name, State, Serial Number, Node ID, IP Address, and Config Source. The table has one row for FAD1. A red arrow points to a gear icon in the Config Source column, with the text 'Click here to edit the HA'.

Host Name	State	Serial Number	Node ID	IP Address	Config Source
FAD1	Standalone	FADV040000146261	0	169.254.16.97	N/A

Configure the required options:

High Availability Setting

Basic

Synchronization

Advanced

Cluster Mode

Standalone

Active-Passive

Active-Active

Active-Active-VRRP

Group Name

group1

Group ID

14

Default: 0 Range: 0-31

Config Priority

40

Default: 100 Range: 0-255

Local Node ID

0

Default: 0 Range: 0-7

Monitor Interface

* port2

* port3

* port4

* port5

Save

Cancel

High Availability Setting

Group ID

14

Default: 0 Range: 0-31

Config Priority

40

Default: 100 Range: 0-255

Local Node ID

0

Default: 0 Range: 0-7

Monitor Interface

✖ port2

✖ port3

✖ port4

✖ port5

Heartbeat Interface

✖ port6

✖ port7

Data Interface

Node List

☒ 0

☒ 1

☐ 2

☐ 3

☐ 4

☐ 5


☐ 6

☐ 7

Save

Cancel

Configure the synchronization options


 **High Availability Setting**

Basic


Synchronization

Advanced


Layer 7 Persistence Synchronization

ON 

Layer 4 Persistence Synchronization

ON 

Layer 4 Connection Synchronization

ON 

Save

Cancel

Configure the advanced options

High Availability Setting

Basic Synchronization **Advanced**

Priority
3
Default: 5 Range: 0-9

Override
ON

Heartbeat Interval
2
Default: 2 Range: 1-20 intervals (100 milliseconds per interval)

Lost Heartbeat Threshold
6
Default: 6 Range: 1-60 retries

ARP Times
5
Default: 5 Range: 1-60 times

ARP Interval
6
Default: 6 Range: 1-20 seconds

Save Cancel

3) Configure the necessary node-ip-list

Typically, you need to configure the IP address for the HA-AA mode. In this mode, the IP address configuration under interface directly is not working. Only the IP address under “config ha-node-ip-list” can work accordingly. In this example, for the port2, the original IP address:

```
config system interface
  edit "port2"
    set ip 159.3.200.4/16
end
```

This IP address 159.3.200.4 is not working. To make it work, we should do the config like this:

```
config system interface
  edit "port2"
    config ha-node-ip-list
      edit 1
        set ip 159.3.200.4/16
        set node 0
        set allowaccess https ping ssh snmp http telnet
      next
      edit 2
        set ip 159.3.200.5/16
        set node 1
        set allowaccess https ping ssh snmp http telnet
      next
    end
  end
```

```
    next
end
```

Then FAD1's port2 uses "159.3.200.4", while FAD2's port2 uses "159.3.200.5".

Deploy HA-VRRP mode

1) Plan the HA deployment

The condition to make sure negotiation successfully:

- All the HA devices use the same heartbeat ports and data ports.
- All the HA devices have same group-id
- All the HA devices have different local-node-id

How the traffic-group-Primary is elected in HA-VRRP mode (Primary and Secondary is elected by traffic-groups):

Preempt enabled:

work state > failover-order > uptime

Preempt disabled:

work state > uptime > failover-order

- Currently, the work state is only impacted by the remote-ip check, if the device contains remote-ip check down, then it deems as down for work state, if one device contains down, while the other doesn't contain, then the one doesn't contain wins. If all the devices contain remote-ip check down, then compare the next condition.
- Failover-order is the option of HA configs. It specifies the alternative device failover order by local-node-id.
- Uptime is the HA device uptime, the more the better.

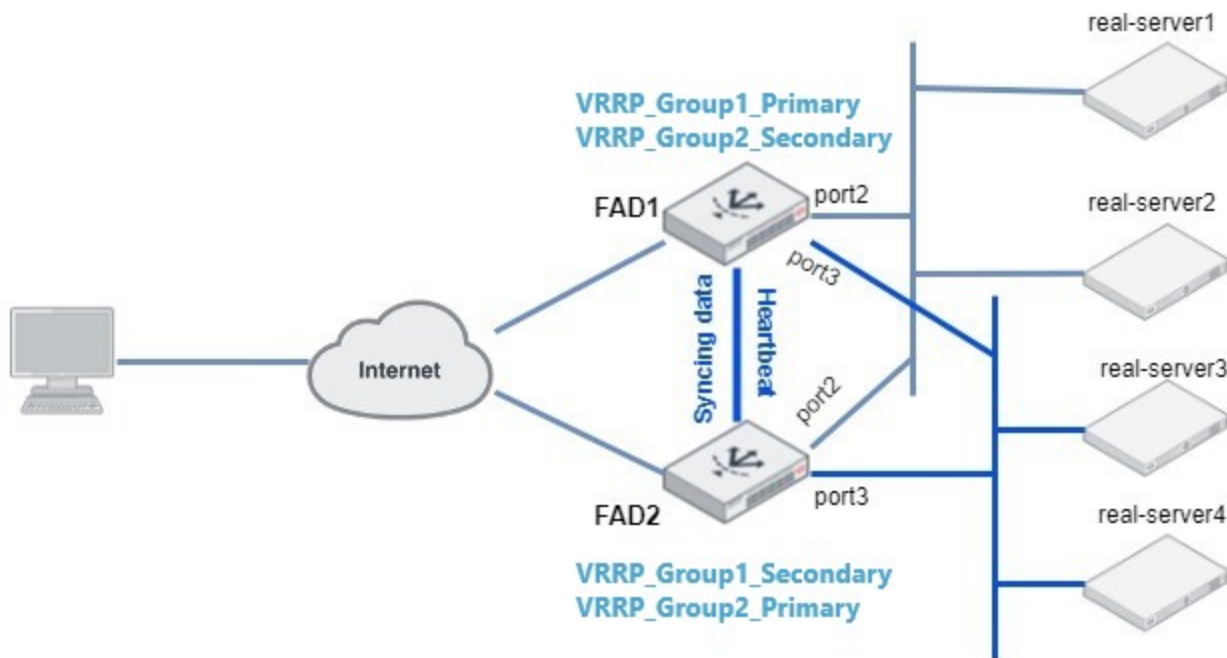
How the config-Primary is elected (This is same in 3 modes):

config-priority > SN

- Config-priority is the value specified in HA config, the device with lower config-priority value will be the config-Primary.
- SN means the serial number, the device with higher SN will be the config-Primary.

Before we get started to set up the HA-VRRP mode, we have to divide the real-servers into groups, typically the real-servers should be divided into 2 groups. In this example, the two groups are VRRP_Group1 and VRRP_Group2.

HA VRRP Mode



Please refer to the chart above; we will make two virtual-servers: VS1, VS2. VS1 belongs to VRRP_Group1, VS2 belongs to VRRP_Group2. The real-server1 and real-server2 belong to VS1, real-server3 and real-server4 belong to VS2. Then all the traffic to VS1 will be handled by FAD1, all the traffic to VS2 will be handled by FAD2. If one of the FortiADC is failing, the other device will take over the traffic. Port2 belongs to VRRP_Group1, port3 belongs to VRRP_Group2.

In this example, we are going to make the HA VRRP config like:

FAD1:

```
config system ha
  set mode active-active-vrrp
  set hbdev port4 port5
  set group-id 15
  set local-node-id 0
  set group-name grp2
  set config-priority 20
  set override enable
  set l7-persistence-pickup enable
  set l4-persistence-pickup enable
  set l4-session-pickup enable
end
```

FAD2:

```
config system ha
  set mode active-active-vrrp
  set hbdev port4 port5
  set group-id 15
  set local-node-id 1
  set group-name grp2
```

```
set config-priority 100
set override enable
set l7-persistence-pickup enable
set l4-persistence-pickup enable
set l4-session-pickup enable
end
```

2) Configure the HA VRRP basic options

In this example, we are going to make FAD1 the config-Primary, FAD2 the config-Secondary. In VRRP mode, each interface has its own IP address, so you can configure the HA-VRRP basic from Web-UI.

The following example shows the FAD1 configuration, the FAD2 is similar.

Navigate to "System->High Availability" page:

Edit the HA node:

High Availability Setting

Basic Synchronization Advanced

Cluster Mode

Standalone Active-Passive Active-Active **Active-Active-VRRP**

Group Name

grp2

Group ID

15

Default: 0 Range: 0-31

Config Priority

20

Default: 100 Range: 0-255

Local Node ID


0

Default: 0 Range: 0-7

Heartbeat Interface

port4 port5

Save Cancel


 High Availability Setting

Basic


Synchronization

Advanced


Layer 7 Persistence Synchronization

ON 

Layer 4 Persistence Synchronization

ON 

Layer 4 Connection Synchronization

ON 

Save

Cancel

High Availability Setting

Basic

Synchronization

Advanced

Priority

Default: 5 Range: 0-9

Override

☒ ON

Heartbeat Interval

Default: 2 Range: 1-20 intervals (100 milliseconds per interval)

Lost Heartbeat Threshold

Default: 6 Range: 1-60 retries

ARP Times

Default: 5 Range: 1-60 times

ARP Interval

Default: 6 Range: 1-20 seconds

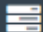
Save

Cancel

3) Configure the needed VRRP groups

Once the two devices established the HA VRRP relationship, then the configuration changes happening on any HA nodes can be synced to the other nodes. So in this example, you can just create the VRRP groups on one of the nodes. Here we put FAD1:

Navigate to System->Traffic Group, add new member

 **Traffic Group**


Traffic Group Name

Preempt
☒ ON

Remote IP Monitor
☐ OFF

Failover Order

Example: 0 1 2 3 4 5 6 7,Range:0-7

 **Traffic Group**

Traffic Group Name

Preempt
☒ ON

Remote IP Monitor
☐ OFF

Failover Order

Example: 0 1 2 3 4 5 6 7,Range:0-7

The equivalent configuration:

```
config system traffic-group
edit "VRRP_Group1"
set failover-order 0 1
set preempt enable
next
edit "VRRP_Group2"
set failover-order 1 0
set preempt enable
next
end
```

4) Assign interface, virtual-server and other resources to the VRRP group

By default, all the interfaces, virtual-servers and other resources are in the traffic-group “default”. We recommend assign the resources to the custom traffic-group.

Navigate to Networking->Interface, edit the interface:

Interface

Name: port2

Status: Up Down

Allow Access: ☒ HTTPS ☒ Ping ☐ SSH ☐ SNMP ☐ HTTP ☐ Telnet

Type: Physical

Traffic Group: VRRP_Group1

Virtual Domain: root

Mode: Static PPPoE DHCP

Floating: ON

Floating IP: 159.9.200.200
Example: 192.0.2.1

Mode Specifics

IPv4/Netmask: 159.9.200.10/16
Example: 192.0.2.5/24

IPv6/Netmask: ::/0
Example: 2001:0db8:85a3::8a2e:0370:7334/64

Secondary IP Address:

Save Cancel

Remember, the float-ip only works on the traffic-group Primary. In this example, the port2 belongs to VRRP_Group1, and FAD1 is currently the Primary of VRRP_Group1, so “159.9.200.200” is only working on FAD1 currently. If FAD1 is failing, then FAD2 will take over the Primary of VRRP_Group1, then the “159.9.200.200” will work on FAD2.

Navigate to Server Load Balance->Virtual Server, edit the interface, set the VS1 to VRRP_Group1, VS2 to VRRP_Group2.

Virtual Server

Basic

General

Monitoring

Name

VS1

Type

Layer 7

Layer 4

Layer 2

Status

Disable

Enable

Maintain

Address Type

IPv4

IPv6

Traffic Group

VRRP_Group1

Specifics

Schedule Pool

OFF

Content Routing

OFF

Packet Forwarding Method

DNAT

Save

Cancel

HA troubleshooting

HA management interface

For HA-AP mode, you are not able to access the Secondary device directly if you didn't enable the management interface. You can only access Secondary CLI from Primary via executing command: "execute ha manage 0" in this scenario. So we recommend you enable the HA management interface for both nodes.

Reminder: Please use "mgmt.-interface" under "config system ha" instead of the old dedicate interface under "config system interface", due to the old dedicate interface has some limitations.

In most cases, you could configure the manage-interface with the IP address same subnets with original port1 or mgmt, so it can be conflict. You'd better have the console control, and then clear the old management IP address on the old interface (typical port1 or mgmt), then set it under "config system ha".

```
FAD2 # config system ha
FAD2 (ha) # set mgmt-status enable
FAD2 (ha) # set mgmt-interface port1
FAD2 (ha) # set mgmt-ip 10.106.188.42/23
FAD2 (ha) # set mgmt-ip-allowaccess http https ping snmp ssh telnet
FAD2 (ha) # end
```

Don't forget to configure the default route accordingly.

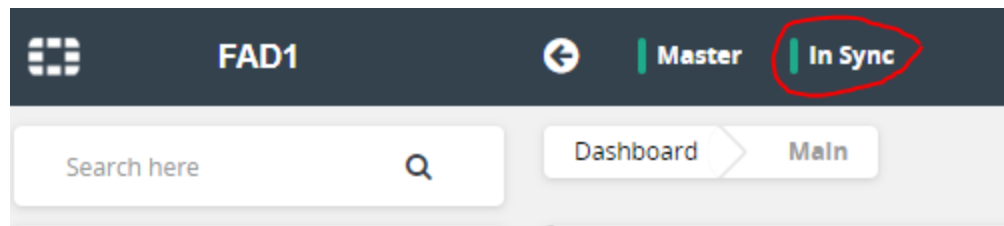
For HA-AA and HA-VRRP mode, you don't have to configure the HA manage-interface, because HA-AA mode uses the IP address of its own local-id, and HA-VRRP mode can have its own interface IP.

For virtualization platform like VMware ESXi, KVM, Hyper-V, please enable the "Promiscuous mode" or "Mac address spoofing" for the management-interface mother interface.

HA config out of sync

Once the HA peers established, all the config can be synced to HA peers by default. There are two kinds of config sync happening, incremental sync and full sync. The incremental sync happens if one of the HA nodes have configuration changes, then the changes will be synced to the HA peers. The full sync happens when the HA daemon restarting triggered, such as the new HA peer joined group.

Normally, you can always see the "In Sync" on the top of the GUI. But if something unknown happened, there could be out of sync happening. In this case, please click the config difference detail at the same position of "In Sync" to see the difference, then you can correct it manually on both devices, or execute the command "execute ha force sync-config" on the correct config side.

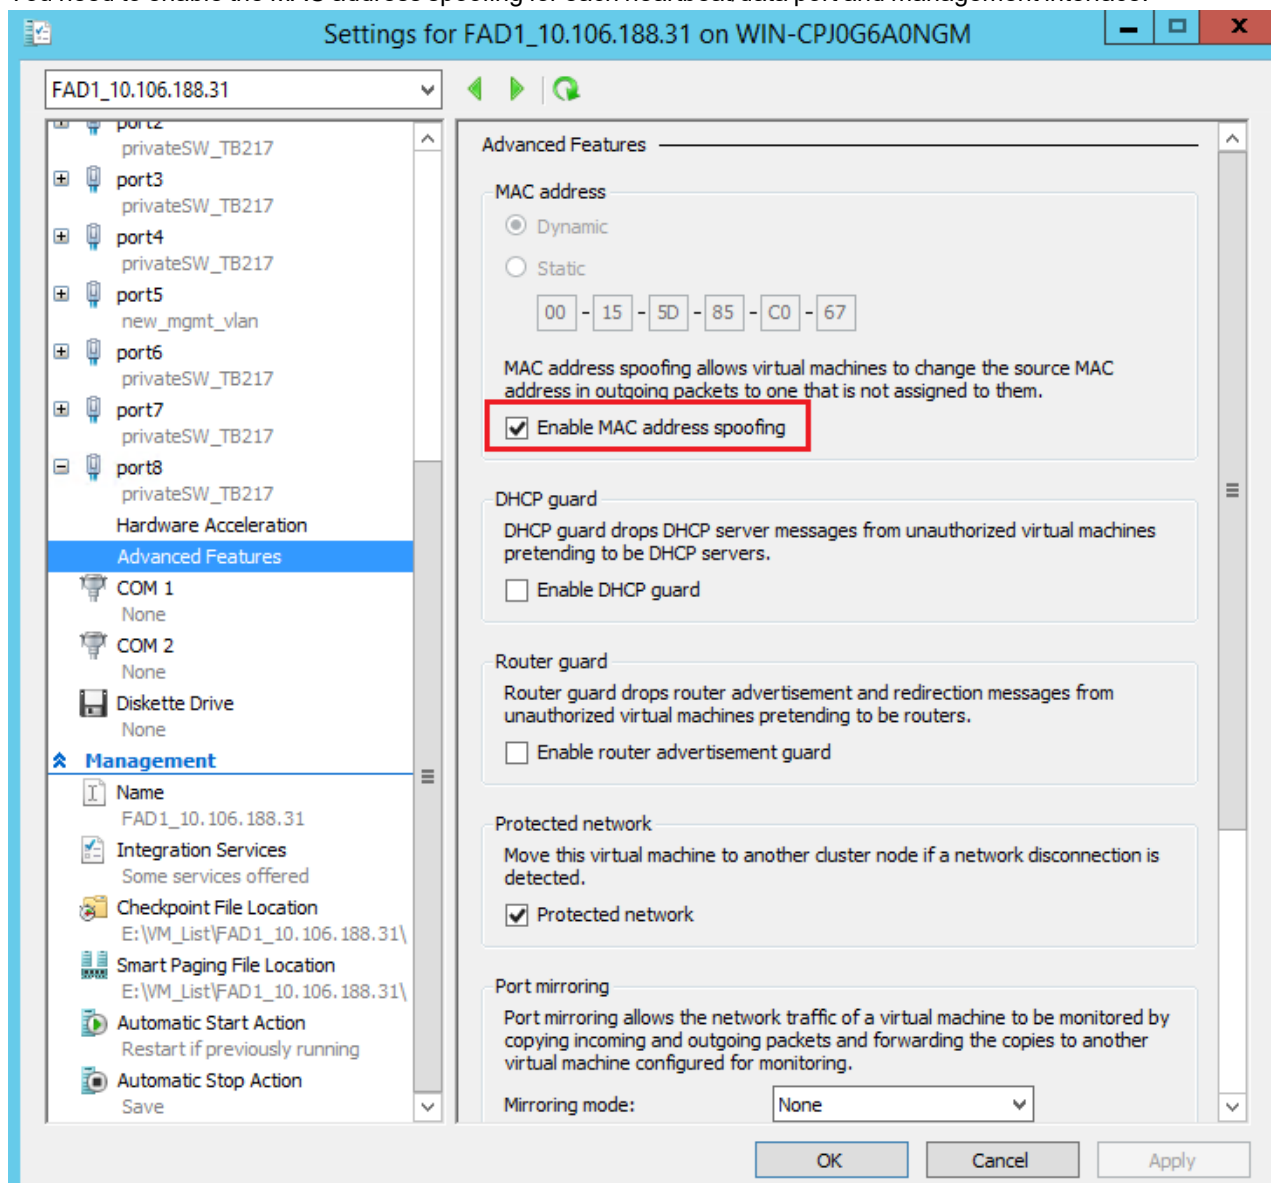


In some scenarios, the above 2 methods can't work. Then you have to backup the correct side full config file, and then restore it to the false HA peer.

HA on Microsoft Hyper-V platform

HA behavior is similar in most platforms except on the Microsoft Hyper-V. The following are some Hyper-V limitations that impact the HA behavior.

1. Unlike other platforms, HA-AP mode on Hyper-V platform utilizes the real MAC address specified by Hyper-V, while other platforms would use the virtual MAC address instead.
2. You need to enable the MAC address spoofing for each heartbeat/data port and management interface.



3. Please assign the individual virtual switch for the heartbeat/data ports due to Hyper-V virtual switch implementation. Otherwise, the HA state could be unstable.
4. It is not supported to set Mac address on Hyper-V platform.

HA abnormal state

If somehow you encounter the HA abnormal state, such as dual Primary, long time waiting to sync. Please check the heartbeat/data ports connectivity. Technically, the heartbeat/data ports should be connected directly, or at least in the same VLAN via switches. If they are connected correctly, then you can enable the debug to see the abnormal reason for the state. Please refer to “4. HA Debug” in this guide to see how to use HA debug command. Here we put an example to debug.

Example:

You can enable the “heartbeat” debug option to see if the heartbeat message was received and sent successfully. If all the heartbeat messages are received and sent properly. Then enable the “errors” option to see if there are errors happening, if so, record it, and try to resolve it. If no more found, please try other options according to the table listed in “4. HA Debug” section of this guide.

Upgrade Firmware

Users can upgrade all FortiADC units within the group with one click by enabling the **HA Cluster Upgrade** in the **Upgrade Firmware UI**.

Upgrade Firmware

HA Cluster Upgrade ☒

Firmware File No file chosen

HA debug

Enable HA debug

1. Enable system debug

```
(M) FAD1 # diagnose debug enable
```

2. Switch on the concerned HA debug options

```
(M) FAD1 # diagnose debug ha basic  
ha debug basic enabled
```

```
(M) FAD1 # diagnose debug ha errors  
ha debug errors enabled
```

3. List all the enabled ha debug options

```
(M) FAD1 # diagnose debug ha list  
basic: enabled  
configuration: disabled  
errors: enabled  
file: disabled  
health-check: disabled  
heartbeat: disabled  
layer4: disabled  
layer7: disabled  
message: disabled  
state: disabled  
sync-status: disabled  
upgrade: disabled  
arp: disabled
```

4. Switch off some HA debug options

```
(M) FAD1 # diagnose debug ha errors  
ha debug errors disabled
```

```
(M) FAD1 # diagnose debug ha list  
basic: enabled  
configuration: disabled  
errors: disabled  
file: disabled  
health-check: disabled  
heartbeat: disabled  
layer4: disabled  
layer7: disabled  
message: disabled  
state: disabled  
sync-status: disabled  
upgrade: disabled  
arp: disabled
```

5. Switch on/off all the HA debug options

```
(M) FAD1 # diagnose debug ha all  
enabled all ha debugs
```

```
(M) FAD1 # [10-09 10:00:58] [kernel]Hello pkt: mode 2 group id 14 local_node_id 0 SN  
FADV040000146260 sented
```

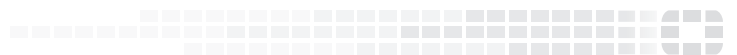
```
[10-09 10:00:58] [kernel]Hello pkt: mode 2 group id 14 local_node_id 0 SN
FADV040000146260 sended
[10-09 10:00:58] [kernel]Hello pkt: mode 2 group id 14 local_node_id 1 SN
FADV040000146261 received
[10-09 10:00:58] [kernel]Hello pkt: mode 2 group id 14 local_node_id 1 SN
FADV040000146261 received
diagnose debug ha all
disabled all ha debugs
```

HA debug options

HA debug option	Meaning
all	Show all the following debug options
arp	Show HA related arp behaviour(especially for GARP)
basic	Show HA basic message.
configuration	Show HA configuration changes sync or full sync.
errors	Show HA errors found.
file	Show HA backend file sync status.
health-check	Show health-check sync status.
heartbeat	Show heartbeat message between HA nodes.
layer4	Show layer4 VS session/persistence table sync status.
layer7	Show layer7 VS persistence table sync status.
list	List all the ha options enabled/disabled status.
message	Show some HA basic message.
state	Show HA state changing log.
sync-status	Show HA sync-status.
updated	Show updated message.
upgrade	Show image upgrading with HA.



FORTINET®



Copyright© 2023 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.