BIG-IP® Virtual Edition and Microsoft Hyper-V: Setup

Version 13.0

Table of Contents

Getting Started with BIG-IP Virtual Edition in Hyper-V

About single NIC and multi-NIC configurations

A typical BIG-IP VE configuration might include four NICs: one for management, one for internal, one for external, and one for high availability.

However, if you want to create a VM for a quick test, you can create a configuration with just one NIC. In this case, BIG-IP VE creates basic networking objects for you.

When BIG-IP VE first boots, it determines the number of active NICs. If BIG-IP VE detects one NIC, then:

- Networking objects (vNIC 1.0, a VLAN named Internal, and an associated self IP address) are created automatically for you.
- The port for the Configuration utility is moved from 443 to 8443.
- High availability (failover) is not supported, but config sync is.
- VLANs must have untagged interfaces.

If BIG-IP VE detects multiple NICs, then you create the networking objects. The port for the Configuration utility remains 443.

You can change the number of NICs after first boot and move from single to multi-NIC and vice versa.

Note: *If there is no DHCP server in your environment and no IP address automatically assigned, then the networking objects will not be created and the port will not be moved.*

Steps to deploy BIG-IP VE

To deploy the BIG-IP® Virtual Edition (VE) system on Hyper-V, you will perform these tasks.

Step	Details
1	Choose the license you want to buy, the BIG-IP VE modules you want, and the throughput you need. See *K14810: Overview of BIG-IP® VE license and throughput limits* on the AskF5™ Knowledge Base (`http://support.f5.com`) for details.
2	Confirm that you are running a hypervisor version that is compatible with a BIG-IP VE release. See *Virtual Edition and Supported Hypervisors Matrix* on `http://support.f5.com` for details.
3	Verify that the host hardware meets the recommended requirements.
4	Download a BIG-IP VE image and deploy it.
5	If you are running a multi-NIC configuration without DHCP, manually assign an IP address for the BIG-IP Config Utility.

After you complete these tasks, you can log in to the BIG-IP VE system and run the Setup utility to perform basic network configuration.

Prerequisites for BIG-IP Virtual Edition on Hyper-V

Host CPU requirements

The host hardware CPU must meet the following requirements.

- The CPU must have 64-bit architecture.
- The CPU must have virtualization support (AMD-V or Intel VT-x) enabled.
- The CPU must support a one-to-one, thread-to-defined virtual CPU ratio, or on single-threading architectures, support at least one core per defined virtual CPU.
- If your CPU supports the Advanced Encryption Standard New Instruction (AES-NI), SSL encryption processing on BIG-IP® VE will be faster. Contact your CPU vendor for details about which CPUs provide AES-NI support.

Host memory requirements

The number of licensed TMM cores determines how much memory the host system requires.

Number of cores	Memory required
1	2 Gb
2	4 Gb
4	8 Gb
8	16 Gb

Virtual machine memory requirements

The guest should have a minimum of 4 GB of RAM for the initial 2 virtual CPUs. For each additional CPU, you should add an additional 2 GB of RAM.

If you license additional modules, you should add memory.

Provisioned memory	Supported modules	Details
4 GB or fewer	Two modules maximum.	AAM can be provisioned as standalone only.
4-8 GB	Three modules maximum.	BIG-IP® DNS does not count toward the module limit. Exception: Application Acceleration Manager™ (AAM®) cannot be provisioned with any other module; AAM is standalone only.
8 GB	Three modules maximum.	BIG-IP DNS does not count toward the module-combination limit.
12 GB or more	All modules.	N/A

Important: *To achieve licensing performance limits, all allocated memory must be reserved.*

Virtual machine storage requirements

The BIG-IP® modules you want to use determine how much storage the guest needs.

Provisioned storage	Supported modules	Details
8 GB	Local Traffic Manager™ (LTM®) module only; no space for LTM upgrades.	You can increase storage if you need to upgrade LTM or provision additional modules.
38 GB	LTM module only; space for installing LTM upgrades.	You can increase storage if you decide to provision additional modules. You can also install another instance of LTM on a separate partition.
127 GB	All modules and space for installing upgrades.	The Application Acceleration Manager™ (AAM®) module requires 20 GB of additional storage dedicated to AAM. For information about configuring the Datastore volume, see *Disk Management for Datastore* on the AskF5™ Knowledge Base (`http://support.f5.com`) for details.

For production environments, virtual disks should be deployed Thick (allocated up front). Thin deployments are acceptable for lab environments.

Note: *To change the disk size after deploying the BIG-IP system, see Increasing disk space for BIG-IP® VE.*

Virtual machine network interfaces

When you deploy BIG-IP® VE, a specific number of virtual network interfaces (vNICs) are available.

For Hyper-V, three virtual NICs are required, at a minimum. More NICs are required if the high availability option is configured.

Each virtual machine can have a maximum of 8 virtual NICs. If you have a legacy virtual NIC, you can use it for management access and still have 8 NICs for dataplane. (The legacy NIC is not supported for the dataplane.) For peak performance, F5® does not recommend using legacy NICs.

Deploying BIG-IP Virtual Edition in Hyper-V

Deploy BIG-IP VE on Hyper-V

To create a BIG-IP ® VE virtual appliance, download an image from F5® and deploy it in your environment.

Important: *Do not change the configuration (CPU, RAM, and network adapters) of the Hyper-V guest environment with settings less powerful than those recommended and described here.*

---

1. In a browser, open the F5 Downloads page (`https://downloads.f5.com`) and log in.
2. On the Downloads Overview page, select **Find a Download**.
3. Under Product Line, select **BIG-IP v12.x/Virtual Edition**.
4. Under Name, select **Virtual-Edition**.
5. If the End User Software License is displayed, read it and then click **I Accept**.
6. Download the BIG-IP VE file package ending with `vhd.zip`.
7. Extract the file from the Zip archive and save it where your VHD files reside on the Hyper-V server.
8. Start Hyper-V Manager, log in to the Hyper-V server, and from the Actions pane, click **New** > **Virtual Machine**.
9. In the **Name** field, type a name for the BIG-IP VE virtual machine and click **Next**.
10. In the **Memory** field, type `4096` and click **Next**.

    Tip: *To increase performance, you can specify a value up to 8192.*

    ---

11. For the **Connection** setting, select **Management** and click **Next**.
12. Map the source network **HA** to the name of a high-availability network in your inventory.
13. Select the **Use an existing virtual hard disk** check box, browse to the location where you saved your VHD file, select the file, open it, and click **Next**.
14. In the Summary page, review the settings and click **Finish**.
    The New Virtual Machine Wizard closes, and the new BIG-IP VE shows in the Virtual Machines list.
15. From the Virtual Machines list, select the new BIG-IP VE.
    The name of the BIG-IP VE appears in the bottom half of the Actions pane.
16. In the lower half of the Actions pane, click **Settings**.
    The Settings window for the selected BIG-IP VE opens.
17. From the Hardware list, select **Processor**, and then change the **Number of logical processors** to `2`, and increase the **Virtual machine reserve (percentage)** to `100`.
18. Click **Add Hardware**, select **Network Adapter** and click **Add**. Repeat this for a total of three adapters (or four for a high-availability configuration).
    a) Click the second Network Adapter, and then from the Network list select **External**.
    b) Click the third Network Adapter, and then from the Network list select **Internal**.
    c) If the BIG-IP VE will be used in a high-availability configuration, click the fourth Network Adapter, and then from the Network list select **HA**.
19. In the Management area, click **Automatic Stop Action** and select **Shut down the guest operating system**.
    This setting ensures that the BIG-IP VE virtual machine restarts with all previously configured virtual hardware and at the current system time.

20. Click **OK** to save your changes and close the Settings window.

Access the BIG-IP VE Configuration utility

If your network has DHCP, an IP address is automatically assigned to BIG-IP® VE during deployment. You can use this address to access the BIG-IP VE Configuration utility or `tmsh` command-line utility.

If no IP address was assigned, you can assign one by using the Configuration utility tool.

1. In the Hyper-V Manager, locate and highlight the virtual machine to which you want to assign the management IP address.
2. In the Actions pane, choose **Connect**.
   The console screen opens. After a few seconds, a login prompt appears.
3. At the password prompt, type `default`.
4. Type `config` and press Enter.
   The F5 Management Port Setup screen opens.
5. Click **OK**.
6. If you want DHCP to automatically assign an address for the management port, select **Yes**. Otherwise, select **No** and follow the instructions for manually assigning an IP address and netmask for the management port.

You can use a hypervisor generic statement, such as `tmsh show sys management-ip` to confirm that the management IP address has been properly assigned.

You can now log into the BIG-IP VE Config utility, and license and provision BIG-IP VE.

After Deploying BIG-IP VE on Hyper-V

Requirements and recommendations for optimum Hyper-V throughput

You can increase throughput of BIG-IP® VE on Hyper-V by following the recommendations described here. F5® has certified throughput with these recommendations; however, they are not required for 1Gbps performance.

Host system recommendations

Optimum settings for the host system include:

- An Intel X520 network interface card (NICs) with two ports, one for the external and one for the internal interface. The NIC used for the management and/or HA interfaces can be a 1G card. For configurations that include an HA VLAN, the virtual switch can either share the NIC used by the management VLAN or use its own NIC.
- Each NIC requires a virtual switch.
- If your are using the Intel X520 NIC, the NIC drivers must be upgraded to a version at least as recent as version 3.8.35.0 NDIS 6.30.
- The driver properties for the 10G NICs must be modified to disable Virtual Machine Queue (VMQ) and Interrupt Moderation. (You can use either the Hyper-V Manager graphic user interface or the PowerShell command line interface to perform this modification.)

Hypervisor recommendations

Optimum settings for the Hyper-V Manager include:

- Disable Non-Uniform Memory Access (NUMA) spanning.
- Disable Virtual Machine Queue (VMQ) for the VE. You should do this even if you've already disabled VMQ for the hypervisor, but especially if you chose not to disable VMQ for the hypervisor NICs.
- Increase the number of licensed TMM cores to 8 and the amount of memory to 16 Gb.

BIG-IP VE considerations

You may also increase your VE performance by reducing the interrupt coalescing threshold for the BIG-IP VE. You can use the following `tmsh` command:

```
tmsh modify sys db scheduler.unicasleeprxlimit.ltm value 16
```

.

Increase disk space for BIG-IP VE

Before proceeding with these steps, use Hyper-V Manager to expand the disk size for the BIG-IP® VE virtual machine and reboot.

Use the BIG-IP VE `tmsh` utility to increase the amount of disk space used by the four BIG-IP VE directories:

- `/config`
- `/shared`
- `/var`

- `/var/log`

Note: *At the time of this release, decreasing the VE disk size is not supported.*

For each directory you want to resize, complete these steps.

1. Use an SSH tool to access the BIG-IP VE `tmsh` utility.
2. From the command line, log in as `root`.
3. List the current size of the directories on your disk so you can determine which ones need to be resized.
   ```
   tmsh show sys disk directory
   ```
4. Expand the size of the directories in which you need additional space.
   ```
   tmsh modify sys disk directory <directory name> new-size <new directory size in 1KB blocks>
   ```
   For example, use `tmsh modify sys disk directory /config new-size 3145740` to increase the size of /config directory to 3145740 1KB blocks (or roughly 3,221,237,760 bytes).
5. To confirm that the command you just submitted is properly scheduled, you can show the new list of directories again.
   ```
   tmsh show sys disk directory
   ```
6. If you change your mind about a submitted size change, you can revoke the size change.
   ```
   tmsh modify sys disk directory /config new-size 0
   ```
   In this example, the size of the `/config` directory is left as is, revoking any scheduled size changes.

   After you submit this sequence of `tmsh` commands, the directory size changes will be scheduled to occur the next time the BIG-IP VE virtual machine (VM) is rebooted.

The next time the VM running BIG-IP VE reboots, the changes are applied.

Enable config sync for BIG-IP VE in Hyper-V

Before you can complete this task:

- Both BIG-IP VEs must be running the same version of BIG-IP VE system software.
- The BIG-IP VEs must not have the same device name. To view the name, use the tmsh command: `list /cm device`. The device name is in the first line that is returned, for example `cm device bigip1 {`. To change the name, use `mv cm device <current_device_name> <new_device_name>`.

Enable config sync communication when you want to automatically or manually synchronize configuration information.

Note: *The following steps apply to single-NIC configuration only. If you have multiple NICs, follow the standard procedures for enabling config sync.*

1. Use an SSH tool to connect to each of the BIG-IP VEs.
2. Ensure you are at the tmsh prompt.
   ```
   tmsh
   ```
3. On each BIG-IP VE, disable functionality that enforces single NIC setup.
   ```
   modify sys db provision.1nicautoconfig value disable
   ```
4. Confirm that the value was set correctly by typing `list sys db provision.1nicautoconfig`.

   The following text is returned: `value "disable"`.
5. On each BIG-IP VE, specify the static private IP address of the BIG-IP VE itself.

```
modify cm device <device_name> configsync-ip <private_ip_address>
```

6. Establish device trust: On one BIG-IP VE, enter the static private IP address of the other BIG-IP VE, along with its user name and password.

   ```
   modify cm trust-domain add-device { ca-device true device-ip
   <peer_ip_address> device-name <peer_device_name> username <peer_username>
   password <peer_password> }
   ```

7. On the same BIG-IP VE as the previous step, create a sync-failover device group with network failover disabled.

   ```
   create cm device-group <device_group_name> devices add { <all-bigip-device-
   names-separated-by-space> } type sync-failover auto-sync enabled network-
   failover disabled
   ```

8. Sync the BIG-IP VE to the other BIG-IP VE.

   ```
   run cm config-sync to-group <device_group_name>
   ```

Change the NIC used for BIG-IP VE management

By default, management traffic goes through the eth0 NIC and data traffic goes through the other available NICs. If you need to use eth0 for data traffic, you can change the NIC that management traffic goes through.

1. Use SSH to connect to BIG-IP VE.
2. If you need to determine which NICs are available, stop TMM by typing `bigstart stop tmm`. Then type `ip addr` to view the list of available NICs.
3. Change the management NIC by typing `tmsh modify sys db provision.managementeth value eth1` where `eth1` is the NIC you want to use for management. You can use any available NIC.
4. Press Enter.
5. Reboot BIG-IP VE by typing `reboot` and pressing Enter.

When BIG-IP VE is running again, you can use eth0 for data.

Note: *If the subnet associated with the management NIC does not have DHCP, you must assign a new IP address by using the BIG-IP Configuration utility tool.*

About routes in a single NIC configuration

If you want to configure a static route that relies on a gateway in the same subnet as the self IP address, you must first disable the setting that enforces single NIC setup:

```
modify sys db provision.1nicautoconfig value disable
```

Confirm that the value is correct by typing `list sys db provision.1nicautoconfig`.

The return value should be `disable`.

If you do not change this value, any time you reboot BIG-IP VE, the manually-configured static route will cause validation errors during `load sys config`.

Change from single NIC to multi-NIC

When you initially boot BIG-IP VE, if it recognized only one NIC, then some network configuration was done automatically. If you want to use more than one NIC, complete the following steps.

1. Use an SSH tool to connect to BIG-IP VE.
2. Return to the default configuration.

   ```
   tmsh load sys config default
   ```

3. Save the changes.

   ```
   tmsh save sys config
   ```

4. Set a database variable so that the number of NICs will be recognized on reboot.

   ```
   setdb provision.1nic enable
   ```

5. Reboot BIG-IP VE.

   ```
   reboot
   ```

When the BIG-IP VE instance is up and running, multiple NICs will be recognized and you can begin configuring BIG-IP VE.

Change from multi-NIC to single NIC

If you have a BIG-IP VE configuration with multiple NICs, you can simplify the configuration to have only one NIC. When you boot the BIG-IP VE and only one NIC is recognized, some networking objects are created automatically for you.

1. Use an SSH tool to connect to BIG-IP VE.
2. Return to the default configuration of BIG-IP VE.

   ```
   tmsh load sys config default
   ```

3. Save the changes.

   ```
   tmsh save sys config
   ```

4. Set a database variable so that the number of NICs will be recognized on reboot.

   ```
   setdb provision.1nic forced_enable
   ```

5. Reboot BIG-IP VE.

   ```
   reboot
   ```

When the BIG-IP VE instance is up and running, it will have a single NIC and related network objects will be created.

Updating BIG-IP VE

Updating BIG-IP VE

You do not need to reinstall BIG-IP® VE in order to install updates. You can use the Software Management tool in the Configuration utility, or you can upgrade the software from the command line.

To update BIG-IP VE, you will:

1. Download the ISO and MD5 files.
2. Install the downloaded files to an inactive boot location.
3. Boot the BIG-IP VE to the new boot location.

Download and import a BIG-IP VE update

To install an update, BIG-IP software needs access to the ISO file. If the update is a hotfix, you need the ISO files for both the base version and the hotfix.

1. In a browser, open the F5® Downloads page (`https://downloads.f5.com`).
2. Download the version's base ISO file and its associated MD5 checksum file.
3. Download the update ISO file and its associated MD5 checksum file.

   Important: *Before you perform the installation, you should test the integrity of the ISO files to verify that you have downloaded clean copies. Use an MD5 verification program to ensure that the downloaded ISO files' checksums match the values in their corresponding MD5 files.*

4. In the BIG-IP VE user interface, on the Main tab, click **System** > **Software Management** > **Image List** > **Import**.
5. Click **Browse** to navigate to the downloaded base level installation file.
6. When the image name appears in the **Software Image** field, click **Import** to begin the operation.

   Important: *Do not navigate away from this screen before the operation is done.*

   When the import is complete, the Import page closes and the downloaded base image displays as an available image.
7. Click the Hotfix List tab.
8. At the right side of the screen, click **Import**.
9. Click **Browse** to navigate to the downloaded hotfix installation file.
10. When the image name appears in the **Software Image** field, click **Import** to begin the operation.

    Important: *Do not navigate away from this screen before the operation is done.*

    When the import is complete, the Import page closes and the downloaded hotfix displays in the list of available images.

You can now install the downloaded update.

Install a BIG-IP VE update

After you download and import the software installation image, you can initiate the installation operation. There are three boot locations on which you can install images on the BIG-IP® system. The process for installing a hotfix or a base version is essentially the same.

1. On the Main tab, click **System** > **Software Management**.
   The Software Management Image List screen opens.
2. In the Available Images area, select the software image you want to install and click **Install**.
   The Install Software Image popup screen opens.
3. Select the disk you want to install the image on, and then type or select a volume name, and click **Install**.

   The upgrade process installs the software on the inactive disk location that you specify. This process usually takes between three and ten minutes.

   Tip: *If there is a problem during installation, you can use log messages to troubleshoot a solution. The system stores the installation log file as* `/var/log/liveinstall.log`.

   The software image is installed.

When the installation operation is complete, you can safely reboot the newly installed volume or partition.

Reboot after a BIG-IP VE update

When the installation operation is complete, you can safely reboot into the newly installed volume or partition.

1. On the Main tab, click **System** > **Software Management**.
   The Software Management Image List screen opens.
2. On the menu bar, click **Boot Locations**.
   The Boot Locations screen opens.
3. In the Boot Location column, click the link representing the boot location you want to activate.
   The properties screen for the boot location opens.
4. Click **Activate**.
   A confirmation screen opens.
5. Click **OK** to initiate the reboot operation.
   The system presents progress messages during the restart operation.

When the BIG-IP® VE system reboot is complete, the system presents the login screen. To configure the system, log in using an account that has administrative permissions.

Legal Notices

Legal notices

Publication Date

This document was published on August 13, 2018.

Publication Number

MAN-0349-08

Copyright

Trademarks

For a current list of F5 trademarks and service marks, see *http://www.f5.com/about/guidelines-policies/trademarks*.

All other product and company names herein may be trademarks of their respective owners.

Patents

This product may be protected by one or more patents indicated at: *https://f5.com/about-us/policies/patents*.

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

Index