

BIG-IP[®] Virtual Edition and Linux KVM: Setup

Version 13.1



Table of Contents

Getting Started with BIG-IP Virtual Edition in KVM.....	5
About single NIC and multi-NIC configurations.....	5
Steps to deploy BIG-IP VE.....	5
Prerequisites for BIG-IP Virtual Edition on KVM.....	7
Host CPU requirements.....	7
Host memory requirements.....	7
Configure SR-IOV on the hypervisor.....	7
Virtual machine memory requirements.....	7
Virtual machine storage requirements.....	8
Virtual machine network interfaces.....	8
Deploying BIG-IP Virtual Edition in KVM.....	9
Deploy BIG-IP VE on KVM.....	9
Use BIG-IP Configuration utility tool to set management IP address.....	10
After Deploying BIG-IP VE on KVM.....	11
About licensing BIG-IP VE.....	11
Configure TSO support.....	11
Turn off LRO or GRO.....	11
Increase performance with multiple queues.....	11
About NIC teaming.....	12
Configure NIC teaming.....	12
Increase disk space for BIG-IP VE.....	12
Change the NIC used for BIG-IP VE management.....	13
About routes in a single NIC configuration.....	14
Change from single NIC to multi-NIC.....	14
Change from multi-NIC to single NIC.....	14
Updating BIG-IP VE.....	15
Updating BIG-IP VE.....	15
Download and import a BIG-IP VE update.....	15
Install a BIG-IP VE update.....	15
Reboot after a BIG-IP VE update.....	16
Legal Notices.....	17
Legal notices.....	17

Getting Started with BIG-IP Virtual Edition in KVM

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 KVM, 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 <i>K14810: Overview of BIG-IP[®] VE license and throughput limits</i> 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 <i>Virtual Edition and Supported Hypervisors Matrix</i> on http://support.f5.com for details.
3	Verify that the host hardware meets the recommended requirements.
4	If you plan to use SR-IOV, enable it on the hypervisor.
5	Download a BIG-IP VE image and deploy it.
6	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 KVM

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

Configure SR-IOV on the hypervisor

To increase performance, you can enable Single Root I/O Virtualization (SR-IOV). You need an SR-IOV-compatible network interface card (NIC) installed and the SR-IOV BIOS must be enabled.

See the KVM documentation for details.

To complete SR-IOV configuration, after you deploy BIG-IP® VE, you must add three PCI device NICs and map them to your networks.

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.

Provisioned memory	Supported modules	Details
		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 amount of storage you need depends on the BIG-IP modules you want to use, and whether or not you intend to upgrade.

Note: These values apply to BIG-IP VE 13.1.0.2 and later.

Provisioned storage	Supported modules	Details
9 GB (LTM_1SLOT)	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.
40 GB (LTM)	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.
60 GB (ALL_1SLOT)	All modules except Secure Web Gateway (SWG); no space for installing upgrades.	The Application Acceleration Manager™ (AAM®) module requires 20 GB of additional storage dedicated to AAM. If you are not using AAM, you can remove the datastore disk before starting the VM.
82 GB (ALL)	All modules except SWG and space for installing upgrades.	The Application Acceleration Manager™ (AAM®) module requires 20 GB of additional storage dedicated to AAM. If you are not using AAM, you can remove the datastore disk before starting the VM.

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 *Increase 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. Each virtual machine can have a maximum of 28 NICs.

Deploying BIG-IP Virtual Edition in KVM

Deploy BIG-IP VE on KVM

You can create an instance of the BIG-IP[®] system that runs as a virtual machine on the host system.

***Important:** Do not change the configuration (CPU, RAM, and network adapters) of the KVM 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 v13.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 `qcow2.zip`.
7. Extract the file from the Zip archive and save it where your `qcow2` files reside on the KVM server.
8. Use VNC to access the KVM server, and then start Virt Manager.
9. Right click **localhost (QEMU)**, and from the popup menu, select **New**.
The Create a new virtual machine, Step 1 of 4 dialog box opens.
10. In the **Name** field, type a name for the connection.
11. Select **import existing disk image** as the method for installing the operating system, and click **Forward**.
12. Type the path to the extracted `qcow` file, or click **Browse** to navigate to the path location; select the file, and then click the **Choose Volume** button to fill in the path.
13. In the **OS type** setting, select **Linux**, for the **Version** setting, select **Red Hat Enterprise Linux 6**, and click **Forward**.
14. In the **Memory (RAM)** field, type the appropriate amount of memory (in megabytes) for your deployment. (For example `4096`, for a 4GB deployment) . From the **CPUs** list, select the number of CPU cores appropriate for your deployment, and click **Forward**.
15. Select **Customize configuration before install**, and click the **Advanced options** arrow.
16. Select the network interface adapter that corresponds to your management IP address, and click **Finish**.
The Virtual Machine configuration dialog box opens.
17. Click **Add Hardware**.
The Add New Virtual Hardware dialog box opens.
18. If SR-IOV is not required, select **Network**.
19. From the **Host device** list, select the network interface adapter for your external network, and from the **Device model** list, select **virtio**. Then click **Finish**.
Do this again for your internal and HA networks.
20. If SR-IOV is required, select **PCI Host Device** and then select the PCI device for to the virtual function mapped to your host device's external VLAN. Then click **Finish**.
Do this again for your host device's internal VLAN and HA VLAN.
21. From the left pane, select **Disk 1**.
22. Click the **Advanced options** button.

23. From the **Disk bus** list, select **Virtio**.
24. From the **Storage format** list, select **qcow2**.
25. Click **Apply**.
26. Click **Begin Installation**.
Virtual Machine Manager creates the virtual machine just as you configured it.

Use BIG-IP Configuration utility tool to set management IP address

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 BIG-IP Config utility tool.

1. Connect to the virtual machine by using the hypervisor's console.
2. At the login prompt, type `root`.
3. At the password prompt, type `default`.
4. Type `config` and press Enter.
The F5 Management Port Setup screen opens.
5. Click **OK**.
6. 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 was set properly.

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

After Deploying BIG-IP VE on KVM

About licensing BIG-IP VE

In order to use BIG-IP VE, you must have a license from F5. Starting with BIG-IP VE version 13.1, you can revoke the license from a virtual machine and re-use it on another virtual machine.

From the Configuration utility, to revoke the license, go to **System > License** and click **Revoke**.

From tmsh, to revoke the license, run the command `tmsh revoke sys license`.

This functionality works for BIG-IP VE BYOL only.

Configure TSO support

The TCP Segmentation Offloading (TSO) feature, which includes support for large receive offload (LRO) and Jumbo Frames, is enabled by default.

Note: You must have the admin user role to enable or disable TSO support.

1. Use an SSH tool to access the BIG-IP[®] VE command line.
2. Log in as `root`.
3. Type `tmsh` to open the `tmsh` utility.
You will see `tmsh` at the prompt once the utility is opened.
4. To determine whether TSO support is currently enabled, use the command:
`show sys db tm.tcpsegmentationoffload`
5. To enable support for TSO, use the command:
`sys db tm.tcpsegmentationoffload enable`
6. To disable support for TSO, use the command:
`sys db tm.tcpsegmentationoffload disable`

Turn off LRO or GRO

Although there are a number of ways to turn off LRO, the most reliable way is to connect to the virtual machine and use the `ethtool` utility.

See the KVM documentation for more details.

Increase performance with multiple queues

You can increase throughput by enabling multiple queues on the host.

1. Shut down the guest.
2. Edit the guest's XML by using a tool like `virsh edit`.
3. In each desired `virtio <interface>` section, add the following element:

```
<driver name='vhost' queues='#' />
```

where # equals the number of queues.

Note: The default number of queues for a KVM guest is 1. The optimum number of queues varies by deployment and depends on throughput needs and total desired guest footprint.

4. Save the file.
5. Start the guest.

About NIC teaming

NIC teaming, also known as *port channels*, *bonding*, or in BIG-IP, *trunking*, is defined by the Link Aggregation standard IEEE 802.1AX-2008 (formerly IEEE 802.3ad). This standard offers both increased bandwidth and link failure redundancy at Layer 2.

In a virtual environment, the same MAC address must be assigned to all teamed/aggregated ports. Either the hypervisor or the guest (in kernel trusted mode) can do this assignment.

SR-IOV uses virtual functions to send packets from two or more separate physical ports to the guest. The hypervisor never sees these packets.

In this case, LACP active monitoring in the guest is not possible because the guest does not receive bridge control packets.

Configure NIC teaming

To increase bandwidth and provide link failure redundancy, enable NIC teaming.

1. On the host, complete the following steps.

- a) Enable SR-IOV.
- b) Expose virtual functions (VFs). For example:

```
echo 4 > /sys/class/net/<device name>/device/sriov_numvfs
```
- c) Create VLANs for the VFs and assign MAC addresses. For example:
 - `ip link set p2p1 vf <VF#> mac "00:01:02:03:00:21"`
 - `ip link set p2p1 vf <VF#> vlan <VLAN ID>`

VF# is 0, 1, and choose an arbitrary MAC address.

Important: Ensure the MAC address is the same for each teamed NIC.

- d) Add the VFs to the guest.
2. Then on BIG-IP VE, complete the following steps.

- a) Create the VLAN.
- b) Create the trunk. For example:

```
tmssh create net <trunkname> interfaces add { 1.1 1.2 }
```
- c) Add the trunk to the VLAN. For example:

```
tmssh create net vlan <vlanname> interfaces add { <trunkname> }
```

For more information on creating trunks in BIG-IP, see the trunking documentation on askf5.com.

Increase disk space for BIG-IP VE

Before proceeding with these steps, use Virt 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.

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.lnicautoconfig value disable
```

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

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.


```
tmsm load sys config default
```
3. Save the changes.


```
tmsm save sys config
```
4. Set a database variable so that the number of NICs will be recognized on reboot.


```
setdb provision.lnic 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.


```
tmsm load sys config default
```
3. Save the changes.


```
tmsm save sys config
```
4. Set a database variable so that the number of NICs will be recognized on reboot.


```
setdb provision.lnic 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 10, 2018.

Publication Number

MAN-0437-08

Copyright

Copyright © 2018, F5 Networks, Inc. All rights reserved.

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described by applicable user licenses. F5 reserves the right to change specifications at any time without notice.

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

A

Automatic Shutdown Action *9*

B

BIG-IP Virtual Edition
updating *15*

C

Configuration utility tool *10*

D

deployment overview *5*
disk size, *See* virtual disk size
downloads
 and importing update *15*
 of ISO base file *15*
 of ISO update file *15*

E

eth0
 changing *13*

G

GRO
 turning off *11*
guest memory *7*

H

hypervisor
 configuring for SR-IOV *7*
 configuring for TSO *11*

I

installation
 rebooting after *16*
installation operation *15*
ISO file
 and location on virtual machine *15*
 downloading *15*
 downloading base file *15*
 downloading update file *15*

K

KVM virtual machine
 creating *9*

L

licensing BIG-IP VE *11*
log file
 and location *15*
 and media speed messages *15*
log in
 after updating *16*
 deploying BIG-IP VE virtual machine *9*
LRO
 turning off *11*

M

management
 accessing different port *13*
memory
 guest *7*
multi-NIC *14*
multiple NIC configuration *5*

N

network adapter
 adding *9*
NIC teaming *12*

P

PEM performance
 optimizing *11*
performance
 increasing *11*

R

reboot operation
 after updating *16*

S

single nic
 routes *14*
single NIC *14*
single NIC configuration *5*
SR-IOV
 configuring hypervisor for *7*
system update procedure *15*

T

task list
 for updating on virtual machine *15*
TSO support
 configuring hypervisor for *11*

Index

U

update

 downloading and importing *15*

 installing *15*

V

VHD file *9*

Virtio *11*

virtual disk size

 extending *12*

 increasing *12*