BIG-IP® Virtual Edition Setup Guide for Linux Xen Project

Version 11.5

Table of Contents

Table of Contents

Getting Started with BIG-IP Virtual Edition

What is BIG-IP Virtual Edition?

BIG-IP® Virtual Edition (VE) is a version of the BIG-IP system that runs as a virtual machine in specifically-supported hypervisors. BIG-IP VE virtualizes a hardware-based BIG-IP system running a VE-compatible version of BIG-IP® software.

Note: The BIG-IP VE product license determines the maximum allowed throughput rate. To view this rate limit, you can display the BIG-IP VE licensing page within the BIG-IP Configuration utility. Lab editions have no guarantee of throughput rate and are not supported for production environments.

About BIG-IP VE compatibility with Xen Project hypervisor products

Each time there is a new release of BIG-IP® Virtual Edition (VE) software, it includes support for additional hypervisor management products. The Virtual Edition and Supported Hypervisors Matrix on the AskF5™ website, `http://support.f5.com`, details which hypervisors are supported for each release.

Important: Hypervisors other than those identified in the matrix are not supported with this BIG-IP version; installation attempts on unsupported platforms might not be successful.

About the hypervisor guest definition requirements

The Xen Project virtual machine guest environment for the BIG-IP® Virtual Edition (VE), at minimum, must include:

*   2 x virtual CPUs
*   4 GB RAM
*   3 x virtual network adapters (minimum); more if configured with the high availability option

    Important: The number of virtual network adapters per virtual machine definition is determined by the hypervisor.

*   1 x 100 GB netfront disk
*   SCSI disk storage; download the image size that provides sufficient space to meet your requirements. An optional secondary disk might also be required as a datastore for specific BIG-IP modules. For information about datastore requirements, refer to the BIG-IP module's documentation.

Note: Refer to Increasing the disk space allotted to the BIG-IP virtual machine for details on changing the disk size after initial download.

Important: You must supply at least the minimum virtual configuration limits to avoid unexpected results.

For production licenses, F5 Networks suggests using the maximum configuration limits for the BIG-IP VE system. For lab editions, required reserves can be less. For each virtual machine, the Xen Project virtual

machine guest environment permits a maximum of 10 network adapters. You can either deploy these as a management port and 9 dataplane ports or a management port, 8 dataplane ports, and an HA port.

There are also some maximum configuration limits to consider for deploying a BIG-IP VE virtual machine, such as:

- CPU reservation can be up to 100 percent of the defined virtual machine hardware. For example, if the hypervisor has a 3 GHz core speed, the reservation of a virtual machine with 2 CPUs can be only 6 GHz or less.
- To achieve licensing performance limits, all allocated RAM must be reserved.
- For production environments, virtual disks should be deployed Thick (allocated up front). Thin deployments are acceptable for lab environments.

Important: There is no longer any limitation on the maximum amount of RAM supported on the hypervisor guest.

Disk space guidelines

The size of the image that you choose to download determines both the number of slots and the number and type of modules that are supported on the VE instance.

Allocated disk space	Supported module combinations	Module specific concerns
7 GB	LTM only on a single slot	You cannot install upgrades or hotfixes to this version.
31 GB	LTM only on two slots.	This option can be extended and upgraded with new versions and hot fix updates. It does not allow installing any modules besides LTM, GTM, or LTM + GTM.
100 GB	Supports all modules. Two slots are supported with potential room to install a third.	This option can be extended and upgraded with new versions and hot fix updates. It allows installing any combination of other modules supported by the current version of BIG-IP VE software.

Guest memory guidelines

The general memory requirement recommendation for BIG-IP® Virtual Edition (VE) is 2 GB per virtual CPU. Additionally, the following memory guidelines may be helpful in setting expectations based on which modules are licensed on VE guests.

Provisioned memory	Supported module combinations	Module specific concerns
12 GB or more	All module combinations are fully supported.	N/A
8 GB	Provisioning more than three modules together is not supported.	GTM™ and Link Controller™ do not count toward the module-combination limit.
More than 4 GB, but less than 8 GB	Provisioning more than three modules together is not supported. (See module-specific concerns relating to AAM).	Application Acceleration Manager™ (AAM) cannot be provisioned with any other module; AAM™ can only be provisioned as Standalone. GTM and Link Controller do not count toward the module-combination limit.
4 GB or less	Provisioning more than two modules together is not supported.	AAM can only be provisioned as Dedicated.

About TCP Segmentation Offloading support

If you want to disable support for TCP Segmentation Offloading (TSO), you must submit a `tmsh` command, because the TSO feature is enabled by default. Note that enabling TSO support also enables support for large receive offload (LRO) and Jumbo Frames.

Configuring a hypervisor for TSO support

You must have the Admin user role to enable or disable TSO support for a hypervisor.

Using the `tmsh` command `sys db`, you can turn TSO support on, off, or check to see whether support is currently enabled.

1. To determine whether TSO support is currently enabled, use the `tmsh list` command.
   ```
   list sys db tm.tcpsegmentationoffload
   ```
2. To enable support for TSO, use the `tmsh enable` command.
   ```
   sys db tm.tcpsegmentationoffload enable
   ```
3. To disable support for TSO, use the `tmsh disable` command.
   ```
   sys db tm.tcpsegmentationoffload disable
   ```

About SR-IOV support

If you want support for SR-IOV, in addition to using the correct hardware and BIOS settings, you must configure hypervisor settings before you set up the guests.

You must have an SR-IOV-compatible network interface card (NIC) installed, and the SR-IOV BIOS enabled before you can configure SR-IOV support.

Refer to the documentation included with your hypervisor operating system for information on support and configuration for SR-IOV.

Deploying BIG-IP Virtual Edition

Host machine requirements and recommendations

To successfully deploy and run the BIG-IP® VE system, the host system must satisfy minimum requirements.

The host system must include:

- CentOS, Debian, Fedora, RHEL, or Ubuntu with the Xen Project package. The *Virtual Edition and Supported Hypervisors Matrix*, published on the AskF5™ web site, http://support.f5.com identifies the Linux versions that are supported for each release, as well as which operating systems provide support for SR-IOV and TSO.
- Virtual Machine Manager®
- Connection to a common NTP source (this is especially important for each host in a redundant system configuration)

If SR-IOV support is required, then additional drivers and configuration steps are required; consult your Linux Xen Project documentation for details.

The hypervisor CPU must meet the following requirements:

- Use 64-bit architecture.
- Have support for virtualization (AMD-V™ or Intel® VT-x) enabled.
- 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 you use an Intel processor, it must be from the Core (or newer) workstation or server family of CPUs.

SSL encryption processing on your VE will be faster if your host CPU supports the Advanced Encryption Standard New Instruction (AES-NI). Contact your CPU vendor for details on which CPUs provide AES-NI support.

The hypervisor memory requirement depends on the number of licensed TMM cores. The table describes these requirements.

Number of Cores	Memory Required
1	2 Gb
2	4 Gb
4	8 Gb
8	16 Gb

About BIG-IP VE Xen Project deployment

To deploy the BIG-IP® Virtual Edition (VE) system on Linux Xen Project, you need to perform these tasks:

- Verify the host machine requirements.
- Deploy an instance of the BIG-IP system as a virtual machine on a host system.

- Power on the BIG-IP VE virtual machine.
- Assign a management IP address to the BIG-IP VE virtual machine.

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

Deploying the BIG-IP VE virtual machine

The primary task in deploying BIG-IP® VE on the open source Xen Project environment is creating and executing a configuration file that sets up most of what you need to get up and running.

Important: Do not modify the configuration of the Xen Project guest environment with settings less powerful than the ones recommended in this document. This includes the settings for the CPU, RAM, and network adapters. Doing so might produce unexpected results.

1. In a browser, open the F5 Downloads page (`https://downloads.f5.com`).
2. Download the BIG-IP VE file package ending with `qcow2.zip`.
3. Extract the file from the Zip archive and save it where your `qcow2` files reside on theXen Project server.
4. Use VNC to access the Xen Project server, and then convert the qcow2 image to the raw format necessary for Xen Project. You can use the following syntax to convert the image.
   ```
   # qemu-img convert <qcow_file_name>.qcow2 <raw_file_name>.raw
   ```
5. Generate a MAC address for the network interface card associated with the management interface.

   Important: Be sure that the MAC address you create starts with the prefix `00:16:3e:`.

   You can use a tool such as MAC Address Generator (`http://www.miniwebtool.com/mac-address-generator/`) to create this address.
6. Use an editor to create a BIG-IP VM definition file that specifies the required parameters for your VM.

   Important: The sample configuration file provided here serves only as an example of the kinds of parameters you need to specify for your virtual machine. The actual file that you create will likely contain different parameters and settings.

   ```
   # vi /etc/xen/<config_file_name>
   ```

   ```
   name = <config_file_name>
   maxmem = 4096
   memory = 4096
   vcpus = 2
   builder = "hvm"
   boot = "c"
   pae = 1
   acpi = 1
   apic = 1
   hpet = 1
   localtime = 0
   on_poweroff = "destroy"
   on_reboot = "restart"
   on_crash = "restart"
   sdl = 0
   vnc = 1
   vncunused = 1
   keymap = "en-us"
   disk = [
   "file:/mnt/xen-bender/bigip/<raw_file_name.raw>,hda,w" ]
   ```

```
                         vif = [
"mac=00:16:<mgmt_interface_mac>,bridge=mgmtbr,script=vif-bridge",

"mac=00:16:3e:<external_interface_mac>,bridge=ext_bridge,script=vif-bridge",


"mac=00:16:3e:<internal_interface_mac>,bridge=int_bridge,script=vif-bridge",]

                         parallel = "none"
                         serial = "pty"
                         #pci = [ '05:10.0', '05:10.1' ]
```

Important: *The last line of the example configuration file contains an optional entry that specifies the IDs for PCI external and internal network interface cards (NIC). This optional entry is required for SR-IOV support. Naturally, if you use this entry, you would omit the external and internal bridges specified in the* vif *section.*

Once you have perfected and saved your configuration file you are ready to create the BIG-IP VM,

7. Run the configuration file using an open source tool such as xm.
   `xm create /etc/xen/<config_file_name>`
   The console should indicate a successful start up by displaying something similar to: `Started domain <config_file_name>(id=444)`

8. Allow some time for the boot-up process; then, you should be able to connect to the BIG-IP console.
   `# xm console <config_file_name>`

Powering on the virtual machine

You power on the virtual machine so that you can begin assigning IP addresses.

There are two default accounts used for initial configuration and setup:

• The root account provides access locally, or using SSH, or using the F5 Configuration utility. The root account password is `default`.
• The admin account provides access through the web interface. The admin account password is `admin`.

You should change passwords for both accounts before bringing a system into production.

Assigning a management IP address to a virtual machine

The virtual machine needs an IP address assigned to its virtual management port.

Tip: *The default configuration for new deployments and installations is for DHCP to acquire the management port IP address.*

---

1. At the password prompt, type `default`.

2. Type `config` and press Enter.
   The F5 Management Port Setup screen opens.

3. Click **OK**.

4. 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.

When assigned, the management IP address appears in the Summary tab of the vSphere™ client. Alternatively, a hypervisor generic statement can be used, such as `tmsh list sys management-ip`

Tip: *F5 Networks highly recommends that you specify a default route for the virtual management port, but it is not required for operating the virtual machine.*

Updating a BIG-IP VE Virtual Machine

About disk space changes

BIG-IP® VE can be deployed in a number of different disk space configurations. Because disk space can be a costly resource, it makes sense to install the smallest disk space configuration that is practical for your operation. If you opt for a smaller disk size, but later determine that you need additional space so you can install upgrades and hot fixes, or to run additional module combinations, you can increase the size of the disk space the BIG-IP VE uses to suit your needs.

Important: At the time of this release, decreasing the VE disk size is not supported. If you decide that you need a smaller disk size, first install a smaller disk size version of BIG-IP VE, and then increase the size to what you want. For example, you might install the 31Gb disk version of BIG-IP VE, but then decide you only need 20Gb. You could install the 7Gb version, and then increase the disk size to 20Gb.

Increasing the disk space allotted to the BIG-IP virtual machine

You can customize the amount of resources available by using a sequence of hypervisor-specific steps to increase the disk space the hypervisor uses, and follow those with a sequence of `tmsh` commands to increase the size of the directories that BIG-IP® VE uses.

1. The process of expanding the virtual disk size of your hypervisor depends on which hypervisor you use. Consult your hypervisor documentation for specific instructions.
2. After you complete the steps for expanding your hypervisor disk size, reboot the BIG-IP VE to see that new disk size.

   Important: Before the BIG-IP VE can use the additional disk space, you need to increase the size of the directories on the disk.

3. Submit the following sequence of `tmsh` commands to schedule expansion of the directories in which you need more room.

   Tip:

   There are four disk directories that can be resized.

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

   a) 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`)
   b) 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).

c) 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`)

d) If you change your mind about a submitted size change, you can revoke that 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. That is, the next time the BIG-IP VE is rebooted, the disk directories are re-sized.

4. Reboot the BIG-IP VE.
During the reboot, the directory size increases finish.

About updates to the BIG-IP VE virtual machine

BIG-IP® VE updates within the same major version are installed in the same manner as updates to BIG-IP software already installed on BIG-IP hardware. You do not need to reinstall BIG-IP VE in the hypervisor guest environment to upgrade your system. To update a BIG-IP VE virtual machine, you can use the Software Management tool in the Configuration utility, or you can upgrade the software from the command line. The update procedure described in this guide uses the Software Management tool.

Downloading and importing 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 before you can successfully import and install a hotfix update.

1. In a browser, open the F5 Downloads page (`https://downloads.f5.com`).
2. Download the version's base ISO file, such as `11.5`, and its associated MD5 checksum file.

> *Tip: The location to which you download ISO files is not significant. Just make sure you can navigate to that location when you perform the import task.*

3. Download the update ISO file, such as `Hotfix-BIGIP-11.5.1-511.0-HF3.iso`, and its associated MD5 checksum file.

> *Important: Before you perform the installation, F5 recommends testing 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 file's checksums match the values in their corresponding MD5 files.*

4. On the Main tab, click **System** > **Software Management** > **Image List** > **Import**.
The Import Software Image screen opens.
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 completes; the system might not import the image successfully. Wait for the operation to complete before continuing with any other work on the BIG-IP VE system.*

The system presents a progress indicator during the operation; 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.
   The Available Images portion of the screen displays any hotfixes you've imported previously.

8. At the right side of the screen, click **Import**.
   The Import Hotfix screen opens.

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 completes; the system might not import the image successfully. Wait for the operation to complete before continuing with any other work on the BIG-IP VE system.*

The system presents a progress indicator during the operation; when the import is complete, the Import page closes and the downloaded hotfix displays in the list of available images.

Installing 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 of the navigation pane, 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.

Rebooting 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 of the navigation pane, 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 December 9, 2016.

Publication Number

MAN-0472-02

Copyright

Trademarks

Patents

Index

Index