BIG-IP® Virtual Edition Setup Guide for VMware ESXi

Version 11.6

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 creates a virtual instance of 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 VMware 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 VMware virtual machine guest environment for the BIG-IP® Virtual Edition (VE), at minimum, must include:

- 2 x virtual CPUs
- 4 GB RAM
- 1 x VMXNET3 virtual network adapter or Flexible virtual network adapter (for management)
- 1 x virtual VMXNET3 virtual network adapter (three are configured in the default deployment for dataplane network access)
- 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 VMware virtual machine guest environment permits a maximum of 10 virtual network adapters (either 10 VMXNET3 with 1 management + 9 dataplane or 1 Flexible management + 9 VMXNET3 dataplane).

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

This table details the capabilities and limitations of the disk space options you can choose.

Provisioned disk space	Capabilities and Limitations	Special Considerations
7 GB	The Local Traffic Manager (LTM®) module is supported, but there is no space available for installing LTM upgrades.	Disk space can be increased if you need upgrade LTM, or decide to provision additional modules.
31 GB	The LTM module is supported. There is also sufficient space available for installing LTM upgrades.	Disk space can be increased if you decide to provision additional modules. You can also install another instance of LTM on a separate partition.
104 GB (OS only) or 124 GB (with Datastore)	All modules and combinations are supported. There is also sufficient space available for installing upgrades.	If you plan to use the Acceleration Module (AM) in addition to other modules, you must add a second 20 GB disk in addition to the 104 GB operating system disk used by the other modules. The 20 GB volume serves as a dedicated Datastore for AM. Do not use this volume for any other purpose. If you need additional space, increase the disk space allotted to this VE. For information on configuring the Datastore volume, refer to *Disk Management for Datastore* published on the AskF5™ web site, `http://support.f5.com`.

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.

Provisioned memory	Supported module combinations	Module specific concerns
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 show` command.
   `show 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.

Configuring a hypervisor for SR-IOV support

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.

From the hypervisor console use `esxcli` (the vSphere command line interface tool) commands to set the system module parameters for `max_vfs`.

1. Check to see what the ixgbe driver settings are currently.
   `esxcli system module parameters list -m ixgbe`
2. Check to see what the ixgbe driver settings are currently.
   In this example, `16,16` is for a 2 port card with 16 virtual functions.
   `esxcli system module parameters set -m ixgbe -p "max_vfs=16,16"`
3. Reboot the hypervisor so that the changes to take effect.
   When you next visit the user interface, the SR-IOV NIC will appear in the Settings area of the Guest as a PCI device.

4. Using the VMware hypervisor user interface, add a PCI device, and then add two virtual functions.

   ```
   05:10.0 | Intel Corporation 82599 Ethernet Controller Virtual Function
   ```

   ```
   05:10.1 | Intel Corporation 82599 Ethernet Controller Virtual Function
   ```

5. Use either the console command line or user interface to configure the VLANs that will serve as pass through devices for the virtual function. For each interface and VLAN combination, specify a name and a value.

   - Name - *pciPassthru0.defaultVlan*
   - Value - `3001`

You can now power on the virtual machine and begin deploying it.

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:

- VMware ESX or ESXi. The *Virtual Edition and Supported Hypervisors Matrix*, published on the AskF5™ web site, `http://support.f5.com` identifies the versions that are supported.
- For SR-IOV support, you need a network interface card that supports SR-IOV; also, make sure that SR-IOV BIOS support is enabled.
- For SR-IOV support, load the `ixgbe` driver and blacklist the `ixgbevf` driver.
- VMware vSphere client
- Connection to a common NTP source (this is especially important for each host in a redundant system configuration)

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 VMware deployment

To deploy the BIG-IP® Virtual Edition (VE) system on VMware ESXi, 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.
- Configure CPU reservation.

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 a BIG-IP VE virtual machine

To create an instance of the BIG-IP® system that runs as a virtual machine on the host system,complete the steps in this procedure.

Important: *Do not modify the configuration of the VMware 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 `scsi.ova`.
3. Start your vSphere Client and log in.
4. From the vSphere Client File menu, choose Deploy OVF Template.
   The Deploy OVF Template wizard starts.
5. In the Source pane, click **Deploy from file or URL**, and, using the **Browse** button, locate the OVA file, open it, and then click **Next**.
   For example: `\MyDocuments\Work\Virtualization\<BIG-IP_OVA_filename>`
   The OVF Template Details pane opens.
6. Verify that the OVF template details are correct, and click **Next**.
   This displays the End-User License Agreement (EULA).
7. Read and accept the license agreement, and click **Next**.
   The Name and Location pane opens.
8. In the **Name** field, type a name for the BIG-IP VE virtual machine, such as: `smith_big-ip_ve`.
9. In the Inventory Location area, select a folder name and click **Next**.
10. From the **Configuration** list, select the number of CPUs and disks required for your system, and then click **Next**.
11. If the host system is controlled by VMware vCenter, the Host Cluster screen opens. Choose the preferred host and click **Next**. Otherwise, proceed to the next step.
12. In the **Datastore** field, type the name of data source your system will use, in the **Available space** field, type in the amount of space your system needs (in Gigabytes), and then click **Next**.
    The Network Mapping dialog box opens.
13. If SR-IOV support is required, skip this step and perform step 14 instead. Map the Source Networks for Management, External, Internal, and HA to the Destination Networks in your inventory.
    a) Map the source network **Management** to the name of the appropriate external network in your inventory.

       An example of a destination external network is **Management**.

    b) Map the source network **Internal** to the name of a destination non-management network in your inventory.

       An example of a destination internal network is **Private Access**.

    c) Map the source network **External** to the name of the appropriate external network in your inventory.

       An example of a destination external network is **Public Access**.

    d) Map the source network **HA** to the name of a high-availability network in your inventory.

An example of a destination internal network is **HA**.

e) When you have all four destination networks correctly mapped, click **Next** and skip the next (SR-IOV only) step.
The Ready to Complete screen opens.

14. (Perform this step only if SR-IOV support is required.) Add PCI device NICs.

a) Delete the existing Source Networks for External, Internal, and HA.

Important: *Be sure to leave the Source Network for the Management NIC.*

b) Edit the settings for the virtual machine to add a PCI device. Map the new device to the name of the device that corresponds to the VLAN associated with your internal subnet.

Assuming your hypervisor setup was performed correctly, there will be 16 virtual functions on each port (05:10.x and 05:11:x) to which you can choose to map your device.

c) Edit the settings for the virtual machine to add a PCI device. Map the new device to the name of the device that corresponds to the VLAN associated with your external subnet.

d) Edit the settings for the virtual machine to add a PCI device. Map the new device to the name of the device that corresponds to the VLAN associated with your internal HA.

e) When you have all four destination networks correctly mapped, click **Next**.
The Ready to Complete screen opens.

15. Verify that all deployment settings are correct, and click **Finish**.

Powering on the virtual machine

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

1. In the main vSphere client window, click the Administration menu.

2. Select the virtual machine that you want to power on.

3. Click the Summary tab, and in the Commands area, click **Power On**.
The status icon changes to indicate that the virtual machine is on. Note that the system will not process traffic until you configure the virtual machine from its command line or through its web interface.

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. From the main vSphere client screen, click the Administration menu.

2. At the <username> 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. 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 show 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.*

Configuring the CPU reservation

Based on selections you made when you deployed the OVA file, a specific amount of memory is reserved for the BIG-IP VE virtual machine.

CPU is not specifically reserved, so to prevent instability on heavily-loaded hosts, you should reserve it manually.

1. In vSphere, edit the properties of the virtual machine.
2. Click the Resources tab.
3. In the Settings area, click **CPU**.
4. In the Resource Allocation section, use the slider to change the reservation.

   The 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.
5. Click **OK**.

Turning off LRO/GRO from the VE guest to optimize PEM performance

Before you can access the VE guest to turn off LRO and GRO, you must have assigned the guest a management IP address.

To optimize performance if you use the virtual machine with the PEM module, you must turn off large receive offload (LRO) and generic receive offload (GRO) for each network interface card (NIC) that is used to pass traffic. You must also use SR-IOV. Although there are a number of ways to turn off LRO, the most reliable way is to connect to the VE guest and use the `ethtool` utility.

1. Use an SSH tool to access the management IP address of the BIG-IP® VE system.
2. From the command line, log in as `root`.
3. Use the `ethtool` to turn off rx-checksumming for the NIC.
   ```
   ethtool -K eth<X> rx off
   ```

   Important: *In this example, substitute the NIC number for <X>.*
4. Use the `ethtool` to turn off LRO for the NIC.

```
ethtool -K eth<X> lro off
```

Important: *In this example, substitute the NIC number for <X>.*

5. Use the `ethtool` to turn off GRO for the NIC.
   ```
   ethtool -K eth<X> gro off
   ```

   Important: *In this example substitute the NIC number for <X>.*

6. Use the `ethtool` to confirm that LRO and GRO are successfully turned off for the NIC.
   ```
   ethtool -k eth<X>
   ```
   In the system response to your command, you should see this info:
   ```
   generic-receive-offload: off
   ```
   ```
   large-receive-offload: off
   ```
   If either of these responses is `on`, your attempt to turn them off was not successful.

   Important: *In this example substitute the NIC number for <X>.*

7. Repeat the previous three steps for each of the NICs that the BIG-IP VE uses to pass traffic.

With LRO and GRO successfully turned off, the performance of the PEM module on the BIG-IP VE system will have better performance and stability.

You can achieve optimum performance (throughput and stability) with the PEM module only if you enable SR-IOV.

Updating a BIG-IP VE Virtual Machine

About VE disk space

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 here 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 have 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, 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.

Upgrading BIG-IP VE from version 10.x to 11.x

When a new version is available, you can upgrade from BIG-IP® VE version 10.x to the BIG-IP VE version 11.x that is available for your hypervisor.

Note: *The BIG-IP VE commands in this procedure are run using the BIG-IP command-line interface within the hypervisor.*

1. Download the base ISO files and MD5 files for versions 10.2.x and version 11 from the F5 Downloads page (`https://downloads.f5.com`), and verify the downloaded ISO files against their corresponding MD5 files.

   Important: *If you are installing a hotfix, you must import the ISO for the base version of the hotfix before you can successfully import and install a hotfix update. For installation, you select only the hotfix image, but the base version of the hotfix must be in place before the hotfix can be applied to BIG-IP VE.*

2. Save a UCS file offline for the BIG-IP VE system you are upgrading.

   Important: *A single configuration file (SCF) will not suffice for this procedure.*

3. Make sure to note the host name for UCS restoration.

   The command string `tmsh show sys global-settings hostname` will show you the active host name.
4. In the VMware vSphere client, power off (shut down) the BIG-IP VE virtual machine.
5. Take a snapshot and back up the BIG-IP VE virtual machine.
6. Add a CD/DVD drive to the virtual machine using either a physical device or ISO file as appropriate.
7. Configure the target virtual machine to boot from the 10.2.x ISO image as the virtual device default node.

   The virtual device node should be IDE, where (0:0) or (1:0) is bootable. The **Connect at power on** setting should remain selected.
8. Remove the existing virtual disk labeled `Hard disk 1`.

   Important: *Select* ***Remove from virtual machine*** *so that the disk is available for recovery purposes.*

9. Add a hard disk.
   a) Select **Create a new virtual disk** and make the disk size 100 GB.

Important: *Other sizes are not supported, and unexpected problems can occur if a size other than 100 GB is used.*

b) Select the Provision type.

For production nodes, a Flat Disk/Thick Provision is preferred to a Thin Provision. For lab editions, a Thin Provision can be used.

c) Select **Create a new virtual disk**.

Use SCSI (0:0) for the Virtual Device node.

d) After the Summary screen, click **Finish**.

At this point the system has:

- Configured a CD/DVD drive to attach when powering on the BIG-IP VE 10.2.x ISO image.
- Removed the legacy 40 GB disk.
- Created a new 100 GB disk.

10. Select BIG-IP VE from the inventory panel and power it on.

11. At the `[VT100]` prompt, press Enter.

12. You are prompted for an unattended default installation; press Enter to continue.

BIG-IP VE version 10.2.2 and later 10.x versions should proceed normally. Versions 10.2.0 and 10.2.1 will return the message: `No configuration provided for hardware type Z100.` For these versions, use the command `diskinit --style volumes` to format the volume, and then use the command `image2disk --nosaveconfig --instslot HD1.1` to install the image.

This will prepare the disk and automatically install the 10.2.x system onto boot volume HD1.1.

13. When the installation is complete for version 10.2.2 and later, the system prompt shows this message: `Remove media, then press [Enter] to reboot`. Press Enter to reboot the machine. For versions 10.2.0 and 10.2.1, type `reboot` and press Enter.

14. Boot into the 10.2.x BIG-IP VE system.

15. Log in to the console and run `config` to configure the management port's IP address.

16. Copy the 10.2.x and 11 ISO files to `/shared/images`.

17. If a hotfix was installed, copy the hotfix ISO file to `/shared/images`, and then perform a live installation of the hotfix onto the system's HD1.2. When finished, reboot into the hotfix image before restoring the UCS.

18. Copy the UCS to `/var/local/ucs/` on the BIG-IP VE.

19. Run the command `tmsh modify sys global-settings hostname xxx.xxx`, where *xxx.xxx* is the previously collected host name.

20. Restore the UCS. For example, if the UCS is named `webfarm1.ucs` and is stored in `/var/local/ucs`, then the restoration command is `tmsh load sys ucs webfarm1`.
At this point, the system is active with the 10.2.x configuration restored.

21. Perform a live installation of BIG-IP VE 11 on the system.

22. When the live installation is complete, halt and power off the BIG-IP VE virtual machine.

23. Within the hypervisor, open the guest settings for the BIG-IP VE.

24. Remove the previously added CD/DVD drive.

25. Adjust the memory reservation from 2 GB to 4 GB to match the memory configuration.

Important: *BIG-IP VE version 11 includes new high-availability features. You should consider adding another VMXNET3 network interface at this point in order to avoid a reboot at a later time.*

26. Power on the BIG-IP VE virtual machine.
The BIG-IP VE virtual machine will boot into version 11.

Deployment Best Practices

Best practices for deploying BIG-IP VE on VMware

When deploying BIG-IP® Virtual Edition (VE) on a VMware host, use these best practices.

Issue	Recommendation
Redundant system configuration	Run the two units of an active-standby pair on separate physical hosts. You can accomplish this in two ways: either manually create a virtual machine peer on each host, or, if you are using VMware Dynamic Resource Scheduler (DRS), create a DRS rule with the option **Separate Virtual Machine** that includes each unit of the BIG-IP® VE redundant pair.
Live migration of BIG-IP VE virtual machines	Perform live migration of BIG-IP VE virtual machines on idle BIG-IP VE virtual machines only. Live migration of BIG-IP VE while the virtual machine is processing traffic could produce unexpected results.
VMware DRS environments	In DRS environments, perform live migration of BIG-IP VE virtual machines (using VMware vMotion) on idle BIG-IP VE virtual machines only. Live migration of BIG-IP VE while the virtual machine is processing traffic could produce unexpected results. Disable automatic migrations by adjusting the DRS Automation Level to **Partially Automated**, **Manual**, or **Disabled** on a per BIG-IP VE basis.
Resource reservations	By default, BIG-IP VE is deployed with a 2000 or 4000 MHz CPU, and 2, 4, or 8 GB of memory reservation. Together, these reservations typically prevent system instability on heavily loaded hosts and are considered minimal. The 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.
Disable hyper-threading on older processors	F5 Networks recommends turning off Hyper-Threading Technology when using host machines with Intel Pentium 4 era processors. Doing so will prevent possible timing issues with BIG-IP VE. ***Important:*** *Production licenses are not supported on Pentium 4 processors.*

Legal Notices

Legal notices

Publication Date

This document was published on July 29, 2016.

Publication Number

MAN-0347-06

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*

Index

T

U

V