



DEPLOYMENT GUIDE

DEPLOYING F5 WITH VMWARE ESX SERVER



vmware®

technology alliance
PARTNER

Table of Contents

Deploying F5 with VMware ESX Server

| | |
|---|-----|
| Prerequisites and configuration notes | 1-1 |
|---|-----|

Configuring the BIG-IP LTM system for VMware ESX

| | |
|---|-----|
| Load Balancing of Virtual Machine Guest Applications | 1-2 |
| Load balancing behavior with DRS | 1-5 |
| Load Balancing behavior with VMware HA | 1-6 |
| Using LTM to improve hardware capacity in a virtual environment | 1-7 |

Configuring the F5 WebAccelerator module with applications running on VMware

| | |
|---|-----|
| Prerequisites and configuration notes | 2-1 |
| Configuration example | 2-1 |
| Configuring the WebAccelerator module | 2-2 |
| Connecting to the BIG-IP LTM device | 2-2 |
| Creating an HTTP Class profile | 2-2 |
| Modifying the Virtual Server to use the Class profile | 2-3 |
| Creating an Application | 2-4 |

Using BIG-IP GTM to provide global site redirection to a secondary data center

| | |
|--|-----|
| Configuring a self IP address on the BIG-IP LTM | 3-2 |
| Creating a Listener on the GTM | 3-2 |
| Creating data centers on the GTM system | 3-3 |
| Creating the monitor | 3-4 |
| Creating Servers for the data center | 3-5 |
| Creating a GTM pool | 3-6 |
| Creating a wide IP on the GTM | 3-8 |
| Configuring the Wide IP as an MX record using ZoneRunner | 3-9 |

Configuring the WANJet device with VMware ESX Servers

| | |
|--|------|
| Common Scenarios | 4-1 |
| Prerequisites and configuration notes | 4-2 |
| Network Topology | 4-2 |
| Pre-deployment tasks | 4-2 |
| Configuring the WANJet devices for ESX devices | 4-5 |
| Configuring the Operational Mode | 4-5 |
| Configuring the Optimization Policy | 4-6 |
| Configuring the Tuning options | 4-7 |
| Modifying the Application QoS settings | 4-8 |
| Finalizing the Deployment | 4-9 |
| Activating the Optimization | 4-9 |
| Additional Recommendations | 4-10 |
| Changes in the WAN environment | 4-10 |
| Performance Charts | 4-10 |
| SMB/CIFS Considerations | 4-10 |



I

Deploying F5 with VMware ESX Server

- Configuring the BIG-IP LTM system for VMware ESX
- Load Balancing of Virtual Machine Guest Applications
- Load balancing behavior with DRS
- Load Balancing behavior with VMware HA
- Using LTM to improve hardware capacity in a virtual environment

Deploying F5 with VMware ESX Server

Welcome to the F5 Deployment Guide on VMware ESX Server. This document provides guidance and configuration procedures for deploying the BIG-IP Local Traffic Manager (LTM), BIG-IP Global Traffic Manager (GTM), WebAccelerator, and WANJet devices with VMware ESX server and the applications running on those devices.

VMware ESX Server, particularly when deployed as part of VMware Virtual Infrastructure, permits a high degree of flexibility for guest operating systems and applications within a datacenter, or between remote data centers. By using BIG-IP LTM and GTM in conjunction with ESX Server, administrators and application architects can improve scalability, ease administrative overhead, and improve end-user experience for applications hosted within a virtualized environment. For additional resources on F5 and VMware, see the [VMware forum on DevCentral](#).

◆ Important

This guide is different than F5's typical Deployment Guides, as the F5 configuration is highly dependent on which applications that are running on the VMware ESX devices. Therefore, most of this document provides general guidance for deploying F5 devices with VMware, as opposed to specific configuration procedures. Refer to the Deployment Guide specific to your application for specific configuration procedures.

This Deployment Guide is broken up into the following chapters:

- *Configuring the BIG-IP LTM system for VMware ESX*, on page 1-2
- *Configuring the F5 WebAccelerator module with applications running on VMware*, on page 2-1
- *Using BIG-IP GTM to provide global site redirection to a secondary data center*, on page 3-1
- *Configuring the WANJet device with VMware ESX Servers*, on page 4-1

Prerequisites and configuration notes

The following are prerequisites for this solution:

- ◆ This Deployment Guide was tested using VMware Virtual Infrastructure 3. Although many of the basic practices outlined in the *Load Balancing of Virtual Machine Guest Applications*, on page 1-2 section of this document are also valid for VMware Server or other VMware virtualization products, most other sections depend on enterprise-level features only found in the Virtual Infrastructure suite.
- ◆ We recommend running BIG-IP LTM version 9.4 or later.
- ◆ Within the context of this deployment guide, **virtual server** will be used to refer to an IP address and port on a BIG-IP LTM which accepts network traffic. On ESX Server, virtualized operating systems will be referred to as **guests**.

Configuring the BIG-IP LTM system for VMware ESX

This section provides general guidance for deploying the BIG-IP LTM system with VMware ESX devices. This section contains the following topics:

- *Load Balancing of Virtual Machine Guest Applications*
- *Load balancing behavior with DRS, on page 1-5*
- *Load Balancing behavior with VMware HA, on page 1-6*
- *Using LTM to improve hardware capacity in a virtual environment, on page 1-7*

Load Balancing of Virtual Machine Guest Applications

In most ways, an application within an ESX guest behaves much like an application running outside of a virtualized environment. Because a BIG-IP LTM directs traffic to a network address, guests that are defined as members in an LTM pool can be located on any number of ESX servers, and distributed among multiple host servers in any manner.

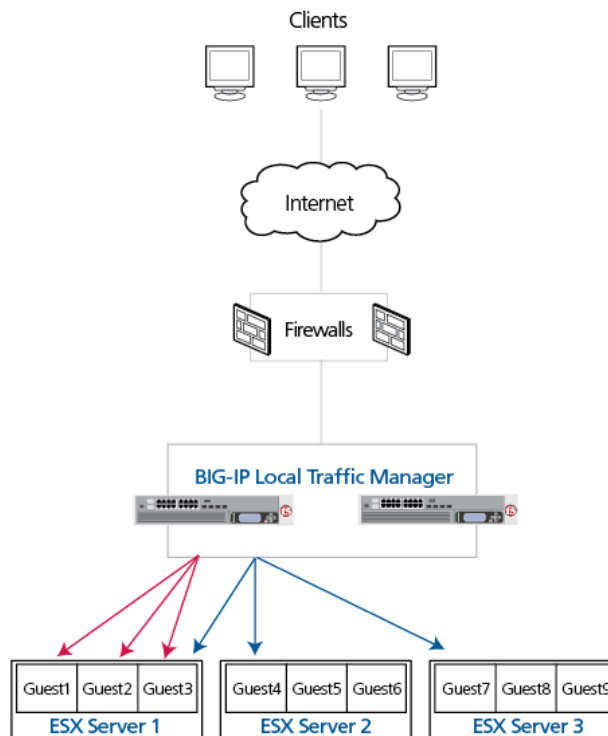


Figure 1.1 Using the BIG-IP LTM to direct traffic to ESX deployments

Considerations for load balancing method

When configuring a BIG-IP LTM system, the IP addresses and Service Ports of the target device or application are added to a load balancing pool. For each BIG-IP LTM pool that contains an ESX device, we recommend choosing one of the following load balancing methods:

◆ **Observed (member)**

The Observed load balancing method allows the BIG-IP LTM to use a combination of logic based on the Least Connections and Fastest load balancing methods to determine the optimal ESX guest to which new traffic should be directed. Since virtualized guests usually co-exist with other applications on the same hardware, this ensures that new traffic is sent to the pool member most able to handle the traffic. For instance, if an ESX server is engaged in heavy disk activity due to events occurring within other guests, and the target guest is therefore unable to process requests in as timely a manner as during normal situations, LTM will dynamically adjust traffic levels to target those guests on other hosts that are better able to process the traffic. The Observed method is particularly useful when ESX hosts may be of dissimilar hardware profiles, or when applications are not evenly distributed throughout an environment.

◆ **Predictive (member)**

The Predictive load balancing method is similar to Observed, except that it also takes into account trending of each pool member. In a highly-dynamic ESX environment, or one that is subject to extreme traffic fluctuations, the Predictive algorithm may help decrease the number of VMotion migrations triggered by VMware DRS technology by directing traffic to guests that not only have the least connections and fastest response times, but those that are likely to remain that way.

To modify the load balancing method of a BIG-IP LTM pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**. The Pool screen opens.
2. From the Pool list, click the name of the applicable pool. The Pool Properties screen opens.
3. On the menu bar, click **Members**.
4. From the **Load Balancing Method** list, select **Observed (member)** or **Predictive (member)** based on the preceding descriptions.
5. Click the **Update** button.




Figure 1.2 Changing the load balancing method of the pool

VMware ESX Server makes it easy to clone existing guests, to suspend guests and resume their operation quickly at a later time, and to fully boot guests that have been shut down. As a result, administrators have the option of quickly adding guests at either predictable intervals, or to meet unexpected traffic spikes. You may want to pre-allocate IP addresses for guests and pre-configure those as nodes and pool members within BIG-IP LTM. No traffic will be sent to those guests unless and until such time as they are brought online, but in that scenario no further configuration of the LTM is required once you make the decision to activate the guests.

Considerations for the health monitor

Health monitors for applications running on ESX should be based on application behavior, not simple methods such as **icmp** or **tcp**. For example, we recommend an advanced health monitor based on the **http** parent that checks for a specific response string from the guest application. This ensures that newly-provisioned, newly-unsuspended, or newly-migrated guests are truly ready to process application traffic correctly.

You can modify an existing health monitor, or create a new one based on the following procedure.

To create an advanced health monitor

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**.
2. Click the **Create** button. The New Monitor screen opens.
3. In the **Name** box, type a name for the Monitor.
In our example, we type **advhttp-monitor**.
4. From the **Type** list, select **http**.
5. In the Configuration section, in the **Interval** and **Timeout** boxes, type an Interval and Timeout. We recommend at least a 1:3 +1 ratio between the interval and the timeout (for example, the default setting has an interval of **5** and an timeout of **16**). In our example, we use a **Interval** of **30** and a **Timeout** of **91**. In the Send String and Receive Rule sections, you can add a Send String and Receive Rule specific to the device being checked.
6. In the **Send String** box, type a string that you expect the target device to return. In our example, we use a Send String of **GET/iisstart.htm**.

If the page you are requesting in the Send String requires authentication, type a user name and password in the appropriate boxes.

7. In the **Receive Rule** box, type what you expect to receive from the Send String. In our example, we expect the Under Construction page to be returned, so we type **[Uu]nder [Cc]onstruction** (see Figure 1.3).

-
- Click the **Finished** button.
The new monitor is added to the Monitor list.

| General Properties | |
|--------------------|-----------------|
| Name | advhttp-monitor |
| Type | HTTP |
| Import Settings | http |

| Configuration: Basic | |
|----------------------|---|
| Interval | 30 seconds |
| Timeout | 91 seconds |
| Send String | GET /iisstart.htm |
| Receive String | [Uu]nder [Cc]onstruction |
| User Name | |
| Password | |
| Reverse | <input type="radio"/> Yes <input checked="" type="radio"/> No |
| Transparent | <input type="radio"/> Yes <input checked="" type="radio"/> No |

Buttons: Cancel Repeat Finished

Figure 1.3 Creating an advanced HTTP monitor

Load balancing behavior with DRS

VMware DRS is a technology for resource balancing, allocation, and management among hosts. Using DRS, an administrator can define threshold limits on each host. When an ESX guest encounters the CPU or memory threshold, DRS moves that guest to a host with more resources. VMware's VMotion technology allows this to occur while the guest is still running.

BIG-IP LTM continues to direct traffic to the guest on the new host; to LTM, this is still the same pool member. By using **Predictive** or **Observed** load balancing methods, traffic is automatically sent to the guest at a level appropriate to new capacities of guest, which is now running on a host that is less constrained.

The BIG-IP LTM offers the additional benefit of mitigating the very brief time in which VMotioned guests are not on the network. A busy web server, for instance, may have several requests interrupted mid-stream. Although many clients will be able to gracefully re-request the data, it's also possible that an application may mis-behave or that content may be delivered

incompletely. By using BIG-IP LTM, which will manage the direct connectivity to the servers on behalf of the clients, you ensure that only valid, correct, and complete data is returned to the client.

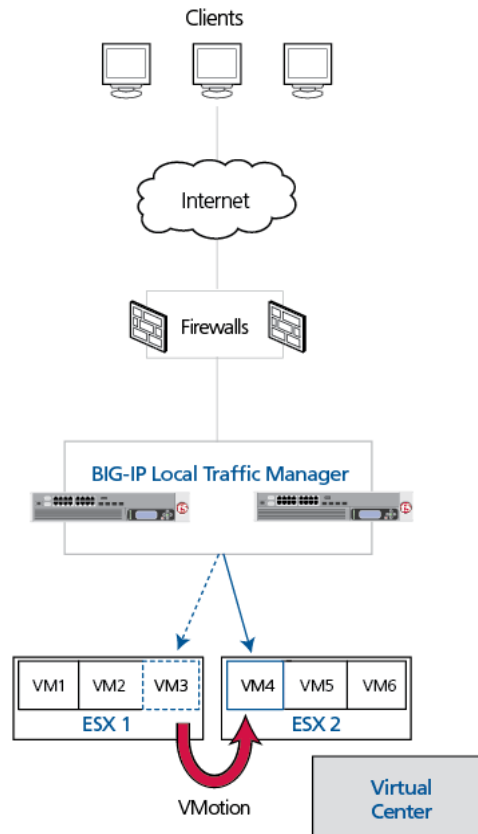


Figure 1.4 Using a BIG-IP LTM to transparently direct appropriate traffic to a VMotioned guest

Load Balancing behavior with VMware HA

VMware HA (High Availability) will restart ESX guests on different hosts, if their original hardware hosts fail. Unlike with VMware DRS, which moves guests to hosts that have more resources available, in a hardware failure situation the total resources available to the resource pool will decrease. As a result, it's likely that guests will restart on hosts with more-constrained capabilities.

BIG-IP LTM will direct traffic to the newly-restarted guests regardless of which host they are assigned. Also, by using **Predictive** or **Observed** load balancing methods, the guests will receive the proportional amount of traffic appropriate to their new hosts.

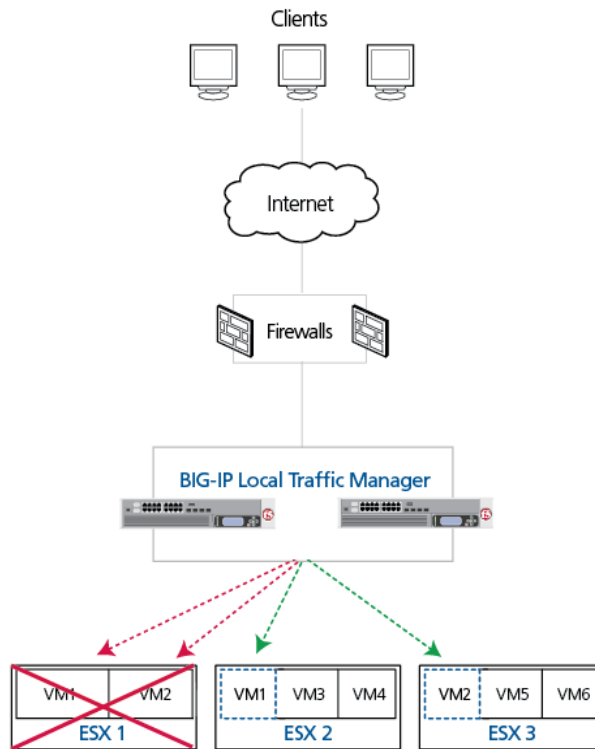


Figure 1.5 Using a BIG-IP LTM to direct traffic to a restarted and correctly responding guest on a different host

Using LTM to improve hardware capacity in a virtual environment

ESX guests share the CPU, disk, and RAM resources of their hosts. By decreasing the per-transaction resources required by each guest, you can dramatically increase the number of virtual machines that can run effectively on any host, while also increasing the effective work that each virtual machine can accomplish.

As an example, BIG-IP's TCP Express feature set ensures optimal network performance for all clients and servers, regardless of operating system and version.

The F5 WebAccelerator (available as a module on the BIG-IP system) can also significantly improve hardware capacity in a virtual environment. See *Configuring the F5 WebAccelerator module with applications running on VMware*, on page 2-1.

Offloading SSL transactions

One of the strengths of the BIG-IP LTM is the ability to terminate HTTPS or other SSL connections, and send traffic to the guests unencrypted. This reduces CPU and memory load on ESX guests by using the dedicated

decryption hardware on the LTM. By terminating SSL/TLS connections at the BIG-IP LTM, you also simplify certificate management, and allow new guest to come online quickly and inexpensively.

To configure the BIG-IP LTM system to offload SSL you need to install a SSL certificate on the BIG-IP LTM and add the certificate and key to a Client SSL profile which is added to the appropriate virtual server. The following procedures describe how to import an SSL certificate into the BIG-IP LTM, how to add the certificate to a profile, and how to modify the virtual server to include the profile.

For information on generating certificates, or using the BIG-IP LTM to generate a request for a new certificate and key from a certificate authority, see the **Managing SSL Traffic** chapter in the *Configuration Guide for Local Traffic Management*.

Importing keys and certificates

Once you have obtained a certificate from a certificate authority, you can import this certificate into the BIG-IP LTM system using the Configuration utility.

To import a key or certificate

1. On the Main tab, expand **Local Traffic**.
2. Click **SSL Certificates**. The list of existing certificates displays.
3. In the upper right corner of the screen, click **Import**.
4. From the **Import Type** list, select the type of import (Certificate or Key).
5. In the **Certificate** (or **Key**) **Name** box, type a unique name for the certificate or key.
6. In the **Certificate** (or **Key**) **Source** box, choose to either upload the file or paste the text.
7. Click **Import**.

If you imported the certificate, repeat this procedure for the key.

Creating a Client SSL profile

The next step is to create a Client SSL profile. This profile contains the SSL certificate and Key information for decrypting the SSL traffic on behalf of the servers.

To create a new Client SSL profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the SSL menu, select **Client**. The Client SSL Profiles screen opens.

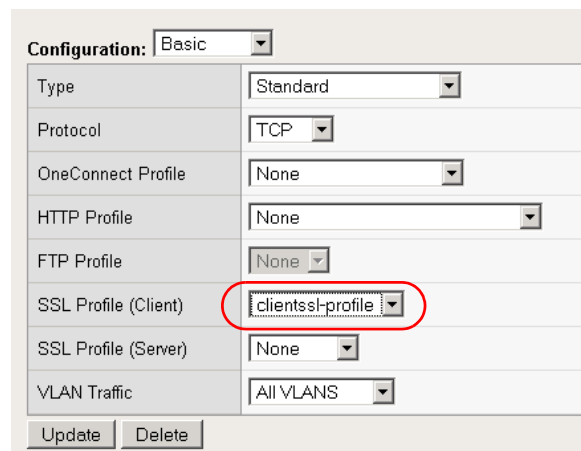
3. In the upper right portion of the screen, click the **Create** button. The New Client SSL Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **clientssl-profile**.
5. In the Configuration section, check the **Certificate** and **Key Custom** boxes.
6. From the **Certificate** list, select the name of the Certificate you imported in the *Importing keys and certificates* section.
7. From the **Key** list, select the key you imported in the *Importing keys and certificates* section.
8. Click the **Finished** button.

Modifying the virtual server to include the Client SSL profile

The final task to enable the BIG-IP LTM to offload SSL is to modify the appropriate virtual server to include the Client SSL profile you just created.

To modify an existing virtual server to use the Client SSL profile

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. From the Virtual Server list, click the virtual server that will be offloading SSL traffic.
3. In the Configuration section, from the **SSL Profile (Client)** list, select the name of the profile you created in *Creating a Client SSL profile*, on page 1-8. In our example, we select **clientssl-profile**.
4. Click the **Update** button.



The screenshot shows a configuration panel for a virtual server. At the top, there is a 'Configuration:' dropdown menu set to 'Basic'. Below this is a table of configuration options:

| | |
|----------------------|-------------------|
| Type | Standard |
| Protocol | TCP |
| OneConnect Profile | None |
| HTTP Profile | None |
| FTP Profile | None |
| SSL Profile (Client) | clientssl-profile |
| SSL Profile (Server) | None |
| VLAN Traffic | All VLANS |

At the bottom of the panel are two buttons: 'Update' and 'Delete'. The 'SSL Profile (Client)' dropdown menu is highlighted with a red circle, indicating the selected profile 'clientssl-profile'.

Figure 1.6 Adding the Client SSL profile to the virtual server

Creating BIG-IP LTM profiles to optimize application transactions

The BIG-IP LTM system uses profiles to enhance your control over managing network traffic, and makes traffic-management tasks easier and more efficient. For applications running on VMware, we recommend using custom HTTP and TCP profiles to optimize the BIG-IP LTM to Guest connections. This allows each guest to perform as efficiently as possible. The optimized HTTP profile makes use of F5's RAM cache and compression engine which speed application transactions.

Although it is possible to use the default profiles, we strongly recommend you create new profiles based on the default parent profiles, even if you do not change any of the settings initially. Creating new profiles allows you to easily modify the profile settings specific to the application, and ensures you do not accidentally overwrite the default profile.

Creating an HTTP profile

The HTTP profile contains numerous configuration options for how the BIG-IP LTM system handles HTTP traffic. In the following example, we leave all settings at their default levels. You can modify any of the profile settings to tune the profile to your application. Although you can use the default profiles, we strongly recommend creating new profiles based off of the parent profile to make

To create a new HTTP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New HTTP Profile screen opens.
3. In the **Name** box, type a name for this profile. In our example, we type **http-optimized**.
4. From the **Parent Profile** list, select **http-wan-optimized-compression-caching**.
5. Check the Custom box for **Content Compression**, and leave **Content List** selected.
6. Modify any of the other settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Creating the WAN optimized TCP profile

The next profile we create is a LAN optimized profile.

To create a new TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **Protocol** menu, click **tcp**.

-
3. In the upper right portion of the screen, click the **Create** button. The New TCP Profile screen opens.
 4. In the **Name** box, type a name for this profile. In our example, we type **optimized-tcp-wan**.
 5. From the **Parent Profile** list, select **tcp-wan-optimized**.
 6. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
 7. Click the **Finished** button.

Creating the LAN optimized TCP profile

The next profile we create is a LAN optimized profile.

To create a new TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **Protocol** menu, click **tcp**.
3. In the upper right portion of the screen, click the **Create** button. The New TCP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **optimized-tcp-lan**.
5. From the **Parent Profile** list, select **tcp-lan-optimized**.
6. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Modifying the virtual server to use the new profiles

The next task is to modify the virtual server to use the new profiles you just created.

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. From the Virtual Server list, click the virtual server that will use the new profiles.
3. From the Configuration list, select **Advanced**. The Advanced configuration options appear.
4. From the **Protocol Profile (Client)** list, select the name of the profile you created in the *Creating the WAN optimized TCP profile* section. In our example, we select **optimized-tcp-wan**.
5. From the **Protocol Profile (Server)** list, select the name of the profile you created in the *Creating the LAN optimized TCP profile* section. In our example, we select **optimized-tcp-lan**.

6. From the HTTP Profile list, select the name of the profile you created in the *Creating an HTTP profile* section. In our example, we select **http-optimized**.
7. Click the **Update** button.

This concludes the BIG-IP LTM system guidance for VMware devices.



2

Deploying the WebAccelerator module with VMware ESX Servers

Configuring the F5 WebAccelerator module with applications running on VMware

In this chapter, we configure the WebAccelerator module for the VMware devices to improve hardware capacity in a virtual environment. The F5 WebAccelerator is an advanced web application delivery solution that provides a series of intelligent technologies designed to overcome problems with browsers, web application platforms and WAN latency issues which impact user performance.

For more information on the F5 WebAccelerator, see www.f5.com/products/big-ip/product-modules/webaccelerator.html.

Prerequisites and configuration notes

The following are prerequisites for this section:

- ◆ We assume that you have already configured the BIG-IP LTM system for directing traffic to the ESX deployment as described in this Deployment Guide.
- ◆ You must have purchased and licensed the WebAccelerator module on the BIG-IP LTM system, version 9.4 or later.
- ◆ If you are using the BIG-IP LTM version 9.4.2 or later, you must have created an HTTP profile on the BIG-IP LTM system that has RAM Cache enabled. In our example we use a parent profile that includes RAM Cache. If you did not create an HTTP profile with RAM Cache enabled, you must create a new HTTP profile, based on a parent profile that uses RAM Cache (we recommend HTTP Acceleration) and associate it with the virtual server. This is only required for BIG-IP LTM version 9.4.2 and later.

Configuration example

Using the configuration in this section, the BIG-IP LTM system with WebAccelerator module is optimally configured to improve hardware capacity for the VMware devices. The BIG-IP LTM with WebAccelerator module both offloads the servers from serving repetitive and duplicate content.

In this configuration, a remote client with WAN latency accesses an ESX device via the WebAccelerator. The user's request is accelerated on repeat visits by the WebAccelerator instructing the browser to use the dynamic or static object that is stored in its local cache. Additionally, dynamic and static objects are cached at the WebAccelerator so that they can be served quickly without requiring the server to re-serve the same objects.

Configuring the WebAccelerator module

Configuring the WebAccelerator module requires creating an HTTP class profile and creating an Application. The WebAccelerator device has a large number of other features and options for fine tuning performance gains, see the *WebAccelerator Administrator Guide* for more information.

Connecting to the BIG-IP LTM device

Use the following procedure to access the BIG-IP LTM system's web-based Configuration utility using a web browser.

To connect to the BIG-IP LTM system using the Configuration utility

1. In a browser, type the following URL:
https://<administrative IP address of the BIG-IP device>
A Security Alert dialog box appears, click **Yes**.
The authorization dialog box appears.
2. Type your user name and password, and click **OK**.
The Welcome screen opens.

Creating an HTTP Class profile

The first procedure is to create an HTTP class profile. When incoming HTTP traffic matches the criteria you specify in the WebAccelerator class, the system diverts the traffic through this class. In the following example, we create a new HTTP class profile, based on the default profile.

To create a new HTTP class profile

1. On the Main tab, expand **WebAccelerator**, and then click **Classes**.
The HTTP Class Profiles screen opens.
2. In the upper right portion of the screen, click the **Create** button.
The New HTTP Class Profile screen opens.
3. In the **Name** box, type a name for this Class. In our example, we type **example-class**.
4. From the Parent Profile list, make sure **httpclass** is selected.
5. In the Configuration section, from the **WebAccelerator** row, make sure **Enabled** is selected.
6. In the Hosts row, from the list select **Match Only**. The Host List options appear.
 - a) In the **Host** box, type the host name that your end users use to access the application on the ESX devices. In our example, we type **http://example-application.f5.com/**.

-
- b) Leave the Entry Type at **Pattern String**.
 - c) Click the **Add** button.
 - d) Repeat these sub-steps for any other host names users might use to access the SAP deployment.
7. The rest of the settings are optional, configure them as applicable for your deployment.
 8. Click the **Finished** button. The new HTTP class is added to the list.

Modifying the Virtual Server to use the Class profile

The next step is to modify the virtual server on the BIG-IP LTM system to use the HTTP Class profile you just created.

To modify the Virtual Server to use the Class profile

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. From the **Virtual Server** list, click the name of the virtual server you created for the application on the VMware devices. In our example, we click **example-http-vs**. The General Properties screen for the Virtual Server opens.
3. On the Menu bar, click **Resources**. The Resources screen for the Virtual Server opens.
4. In the HTTP Class Profiles section, click the **Manage** button.
5. From the **Available** list, select the name of the HTTP Class Profile you created in the preceding procedure, and click the Add (<<) button to move it to the Enabled box. In our example, we select **example-class**.
6. Click the **Update** button. The HTTP Class Profile is now associated with the Virtual Server.

Important

If you are using the BIG-IP LTM version 9.4.2 or later, you must have created an HTTP profile on the BIG-IP LTM system that has RAM Cache enabled. In our example, we use a parent profile that includes RAM Cache. If you did not create an HTTP profile with RAM Cache enabled, you must create a new HTTP profile, based on a parent profile that uses RAM Cache (such as HTTP Acceleration), and modify the virtual server to use this new profile. This is only required for BIG-IP LTM version 9.4.2 and later.

*To create the HTTP profile, use **Creating an HTTP profile**, on page 1-10, selecting the HTTP Acceleration parent profile. You must leave RAM Cache enabled; all other settings are optional. To modify the virtual server, follow*

*Steps 1 and 2 from the preceding procedure to access the virtual server, and then from the HTTP Profile list, select the name of the new profile you just created and click **Update**.*

Creating an Application

The next procedure is to create a WebAccelerator Application. The Application provides key information to the WebAccelerator so that it can handle requests to your application appropriately.

To create a new Application

1. On the Main tab, expand **WebAccelerator**, and then click **Applications**.
The Application screen of the WebAccelerator UI opens in a new window.
2. Click the **New Application** button.
3. In the Application Name box, type a name for your application.
In our example, we type **Example Application**.
4. In the **Description** box, you can optionally type a description for this application.
5. From the **Local Policies** list, select the Policy that best matches the application you are running on the VMware devices. If there is not a predefined policy for your application, you can create a new WebAccelerator policy for your application.
6. In the **Requested Host** box, type the host name that your end users use to access the application. This should be the same host name you used in Step 6a in the preceding procedure. In our example, we type **http://example-application.f5.com/**.
If you have additional host names, click the **Add Host** button and enter the host name(s).
7. Click the **Save** button.

The rest of the configuration options on the WebAccelerator are optional, configure these as applicable for your network. With this base configuration, your end users will notice a marked improvement in performance after their first visit.



3

Deploying the BIG-IP GTM for VMware ESX Server multi-data center deployments

Using BIG-IP GTM to provide global site redirection to a secondary data center

In this chapter, we configure the BIG-IP Global Traffic Manager for multi-data center deployments of VMware. VMware Infrastructure makes hosting duplicate sites or applications in remote data centers an easier and more-manageable task. Because of that, it's even more important than ever to have a global traffic management solution that can direct clients to the correct, functional site in a timely manner.

The BIG-IP Global Traffic Manager module (GTM) can perform all required functions to make this possible. If for instance Site 1 becomes unavailable because its Internet connection is severed, BIG-IP-GTM modifies DNS to direct clients to Site 2 when appropriate -- when replication of VMware images is complete, guests are running, and the application is accepting traffic.

The BIG-IP GTM is available as a module on the BIG-IP system.

For more information on the BIG-IP GTM, see www.f5.com/products/big-ip/product-modules/global-traffic-manager.html




Figure 3.1 Logical configuration example using BIG-IP Local and Global Traffic Managers

Configuring a self IP address on the BIG-IP LTM

The first task in this configuration is to create a unique self IP address on the BIG-IP LTM system for use by the GTM. You need a unique self IP address for each redundant pair of BIG-IP LTM devices in this configuration, so if you have multiple pairs of BIG-IP LTMs you need a unique self IP for each one.

The IP address you choose, and the VLAN to which you assign it, must be accessible by any clients that will be performing DNS queries against the GTM. It may be a private IP address if a Network Address Translation (NAT) device, such as a BIG-IP LTM, a firewall, or a router, is providing a public address and forwarding DNS traffic to the listener.

To create a self IP address

1. On the Main tab, expand **Network**, and then click **Self IPs**.
The Self IP screen opens.
2. Click the **Create** button.
The new Self IP screen opens.
3. In the **IP Address** box, type an IP address in the appropriate VLAN (the VLAN you choose in step 5).
In our example, we type **10.133.20.70**.
4. In the **Netmask** box, type the corresponding subnet mask.
In our example, we type **255.255.255.0**.
5. From the **VLAN** list, select the appropriate VLAN.
6. Click the **Finished** button.
The new self IP address appears in the list.

Creating a Listener on the GTM

The next task is to create a listener on the BIG-IP GTM system. A listener instructs the Global Traffic Manager to listen for network traffic destined for a specific IP address. In our case, this specific IP address is the self IP address on the LTM system we just created.

To create a listener on the GTM system

1. On the Main tab of the navigation pane, expand **Global Traffic** and then click **Listeners**. The main listeners screen opens.
2. Click the **Create** button.
3. In the **Destination** box, type the self IP address you created in *Configuring a self IP address on the BIG-IP LTM*, on page 3-2. In our example, we type **10.133.20.70** (see Figure 3.2).
4. Leave the **VLAN Traffic** list set to **All VLANs**.
5. Click the **Finished** button.

-
- Repeat this procedure for any additional self IP addresses you configured in the *Configuring a self IP address on the BIG-IP LTM* section.




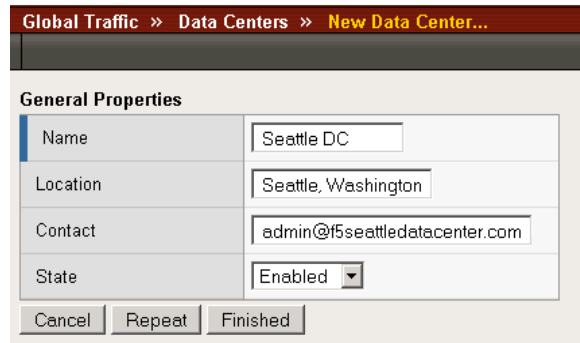
Figure 3.2 Creating a new listener

Creating data centers on the GTM system

The next step is to create data centers on the GTM system for each real-world location that will host globally load balanced ESX devices servers. A data center defines the group of Global Traffic Managers, Local Traffic Managers, host systems, and links that share the same subnet on the network. In our example, we created a Seattle data center and a New York data center.

To create a new Datacenter on the GTM system

- On the Main tab of the navigation pane, expand **Global Traffic** and click **Data Centers**. The main screen for data centers opens.
- Click the **Create** button.
The New Data Center screen opens.
- In the **Name** box, type a name for this datacenter. In our example, we type **Seattle DC**.
- In the **Location** box, type a location that describes the physical location of the data center. In our example, we type **Seattle, Washington**.
- In the **Contact** box, type the name of the person responsible for managing the network at the data center. In our example, we type **admin@f5seattledatacenter.com**.
- Make sure the **State** list remains at **Enabled** (see Figure 3.3).
- Click the **Finished** button.
- Repeat this procedure for each of your data centers. In our example, we repeat the procedure once for our New York data center.



The screenshot shows a web-based configuration window titled "Global Traffic » Data Centers » New Data Center...". Below the title bar is a section labeled "General Properties" containing a form with the following fields:

| | |
|----------|-------------------------------|
| Name | Seattle DC |
| Location | Seattle, Washington |
| Contact | admin@f5seattledatacenter.com |
| State | Enabled |

At the bottom of the form are three buttons: "Cancel", "Repeat", and "Finished".

Figure 3.3 Creating a new GTM data center

Creating the monitor

The next task is to create a monitor on the GTM system. Monitors verify connections on pools and virtual servers and are designed to check the status of a pool or virtual server on an ongoing basis, at a set interval. If a pool or virtual server being checked does not respond within a specified timeout period, or the status of a pool or virtual server indicates that performance is degraded, then the Global Traffic Manager can redirect the traffic to another resource.

In our example, the application running on our VMware devices is an email application, so we create a SMTP monitor. The SMTP monitor issues standard Simple Mail Transport Protocol (SMTP) commands to ensure that the BIG-IP LTM virtual server is available. You can configure a monitor most appropriate for your configuration.

Although it is possible to use the default monitor, we recommend creating a new monitor based off the default monitor, which enables you to configure specific options.

To create a BIG-IP GTM health monitor

1. On the Main tab of the navigation pane, expand **Global Traffic** and then click **Monitors**.
2. Click the **Create** button. The New Monitor screen opens.
3. In the **Name** box, type a name for the monitor. In our example, we type **gtm_smtp**.
4. From the **Type** list, select **SMTP**.
5. Configure the options as applicable for your deployment. In our example, we leave the options at their default levels.
6. Click the **Finished** button.
The new monitor is added to the list.

Creating Servers for the data center

The next task is to create a *GTM Server* for the data centers. A server defines a specific system on the network. In this deployment, the GTM servers are the BIG-IP LTM systems we configured earlier in this guide.

To create a GTM server

1. On the Main tab of the navigation pane, expand **Global Traffic** and click **Servers**.
The main screen for servers opens.
2. Click the **Create** button. The New Server screen opens.
3. In the **Name** box, type a name that identifies the Local Traffic Manager. In our example, we type **Seattle_BIG-IP**.
4. From the Product list, select either **BIG-IP System (Single)** or **BIG-IP System (Redundant)** depending on your configuration. In our example, we select **BIG-IP System (Redundant)**.
5. From the Address List section, in the **Address** box, type the self IP address of the BIG-IP LTM device, and then click the **Add** button. In our example, we type **10.133.20.227**.
6. If you selected BIG-IP System (Redundant) in Step 4, from the Peer Address List section, in the **Address** box, type the self IP address of the redundant BIG-IP LTM device, and then click the **Add** button.

Note: Do not use a floating IP address of the redundant pair. Do not use the administrative interface of the either member of a redundant pair.

7. From the **Data Center** list, select the name of the data center you created in the *Creating data centers on the GTM system* section. In our example, we select **Seattle DC**.
8. In the Health Monitors section, from the Available list, select the name of the monitor you created in the *Creating the monitor* section, and click the Add (<<) button. In our example, we select **gtm_smtp**.
9. In the Resources section, from the Virtual Server Discovery list, choose an option. We recommend **Enabled (No Delete)**. With this option, the GTM will discover all the virtual servers you have configured on the LTM(s) via iControl, and will update, but not delete them.
10. Click the **Finished** button.

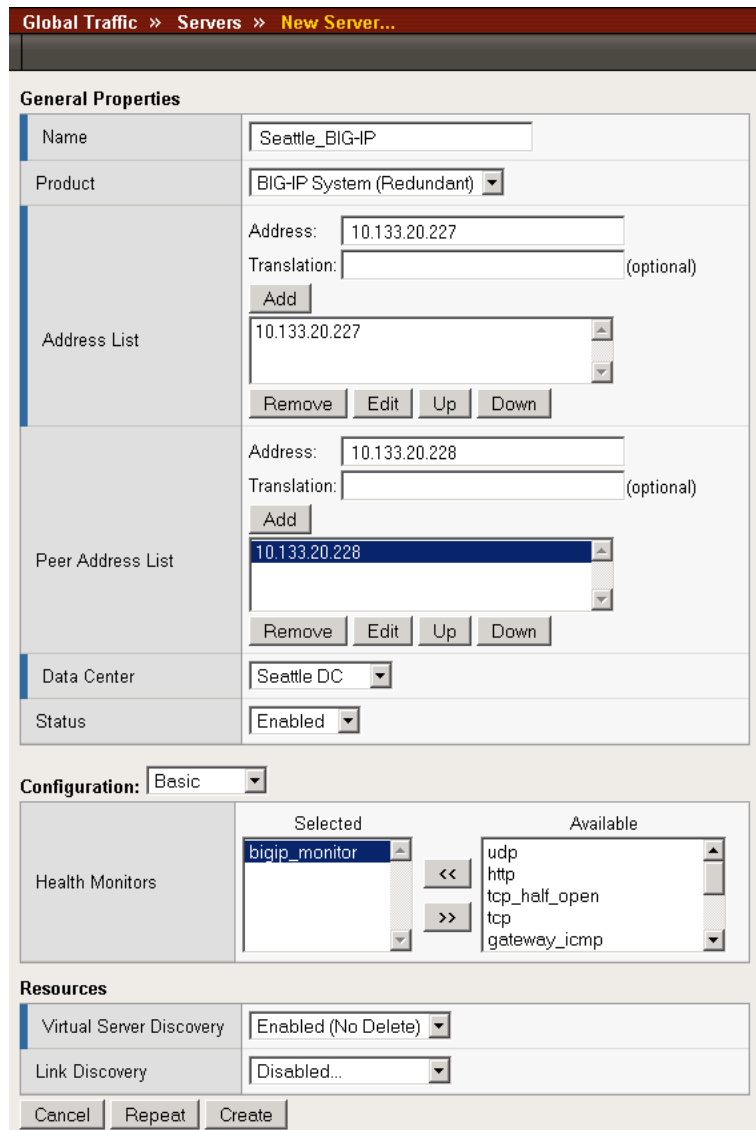


Figure 3.4 Creating a GTM server

Creating a GTM pool

The next task is to create a pool on the GTM device that contains the BIG-IP LTM virtual server.

To create a pool on the GTM

1. On the Main tab of the navigation pane, expand **Global Traffic** and click **Pools** (located under Wide IPs).

2. Click the **Create** button. The New Pool screen opens.
3. In the **Name** box, type a name for the pool. In our example, we type **Seattle_pool**.
4. In the Health Monitors section, from the Available list, select the name of the monitor you created in the *Creating the monitor* section, and click the Add (<<) button. In our example, we select **gtn_smtp**.
5. In the Load Balancing Method section, choose the load balancing methods from the lists appropriate for your configuration. In our example, we select **Global Availability**, **Round Robin**, and **Return to DNS**, in that order.
6. In the Member List section, from the **Virtual Server** list, select the virtual server you created for the ESX devices, and click the **Add** button. You must select the virtual server by IP Address and port number combination. In our example, we select **10.133.20.200:25**. If you have additional virtual servers for the ESX devices configured on the BIG-IP LTM system, repeat this step.
7. Click the **Finished** button.

The screenshot displays the 'New Pool...' configuration window in the BIG-IP LTM interface. The breadcrumb path is 'Global Traffic >> Pools >> New Pool...'. The 'General Properties' section includes a 'Name' field with 'Seattle_pool' and a 'State' dropdown set to 'Enabled'. Below this, the 'Configuration' is set to 'Basic'. The 'Health Monitors' section features two lists: 'Active' (containing 'gtn_smtp') and 'Available' (containing 'udp', 'http', 'tcp_half_open', 'tcp', and 'gateway_icmp'), with left and right arrow buttons between them. The 'Load Balancing Method' section has three dropdowns: 'Preferred' (Global Availability), 'Alternate' (Round Robin), and 'Fallback' (Return to DNS). The 'Fallback IP' field is empty. The 'Members' section includes a 'Virtual Server' dropdown (10.133.16.100:80), a 'Ratio' field (1), an 'Add' button, and a 'Member List' containing '10.133.20.200:25 Ratio(1)'. At the bottom of the window are 'Cancel', 'Repeat', and 'Finished' buttons.

Figure 3.5 Creating a pool on the BIG-IP GTM

Creating a wide IP on the GTM

The next task is to create a wide IP on the GTM system. A *wide IP* is a mapping of a fully-qualified domain name (FQDN) to a set of virtual servers that host the domain's content.

To create a wide IP on the GTM system

1. On the Main tab of the navigation pane, expand **Global Traffic** and click **Wide IPs**.
2. Click the **Create** button. The New Wide IP screen opens.
3. In the **Name** box, type a name for the Wide IP. In our example, we type **mail.example.com**.
4. In our example, we are not using any iRules, so we skip the iRule section. Configure as appropriate for your deployment.
5. In the Pools section, from the **Load Balancing Method** list, select a load balancing method. In our example, we select **Global Availability**. Global Availability instructs the GTM to select the first pool in the wide IP until it becomes unavailable, at which point it selects the next pool until the first pool becomes available again.

In our example, the GTM sends all incoming email to the first-listed pool, **Seattle_pool**. If that pool is unavailable, all incoming email is sent to the next-listed pool, **NewYork_pool**. If you wish to distribute incoming email among multiple pools, select another method, such as Ratio.

Consult the online documentation or the product manual for more details about load balancing methods.

6. From the Pool List section, from the **Pool** list, select the name of the pool you created in the *Creating a GTM pool* section, and then click the **Add** button.
In our example, we select **Seattle_pool**.
Repeat this step for any additional pools. In our example, we repeat one time for the **NewYork_pool**.
7. All other settings are optional, configure as appropriate for your deployment.
8. Click the **Finished** button (see Figure 3.6).

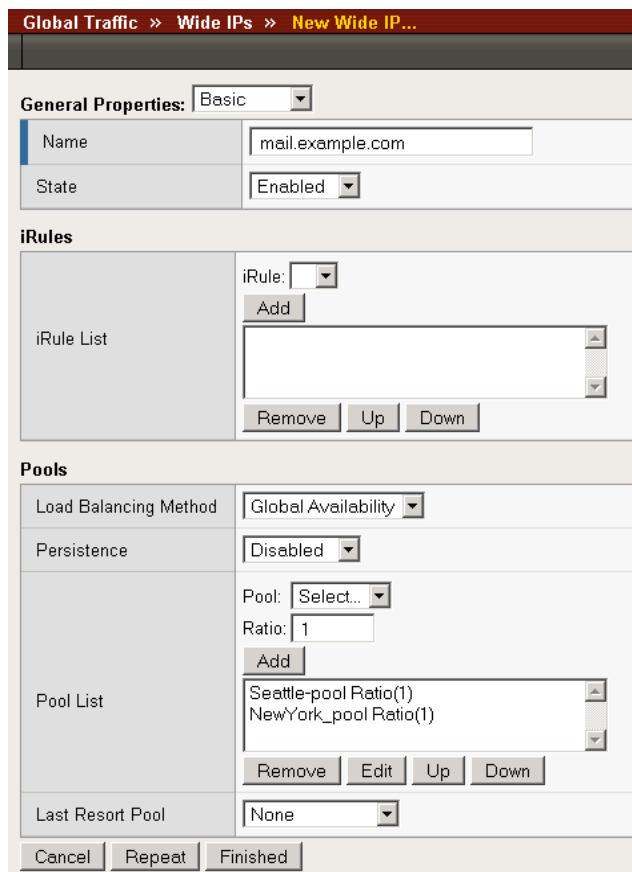


Figure 3.6 Creating a new Wide IP on the GTM system

The next task is to add the newly-created Wide IP as an MX record in your DNS system. If using the GTM as your primary DNS system, this is done through the ZoneRunner utility.

Configuring the Wide IP as an MX record using ZoneRunner

The final task in this configuration is to configure the Wide IP as an MX record in a DNS system. In our example, we are using the GTM system as our primary DNS, and use ZoneRunner to add the Wide IP as an MX record.

The ZoneRunner utility is an advanced feature of the Global Traffic Manager. We highly recommend that you become familiar with the various aspects of BIND and DNS before you use this feature. For in-depth information, we recommend the following resources:

- DNS and BIND, 4th edition, Paul Albitz and Cricket Liu
- The IETF DNS documents, RFC 1034 and RFC 1035

- The Internet Systems Consortium web site,
<http://www.isc.org/index.pl?sw/bind/>

For information on adding the required MX record to other DNS servers, for instance BIND or Microsoft Windows 2007 DNS Service, consult the appropriate product documentation.

To add the Wide IP as an MX record using ZoneRunner

1. On the Main tab of the navigation pane, expand **Global Traffic** and click **ZoneRunner**.
2. Click the **Create** button. The New Resource Record screen opens.
3. From the **View** list, select a view. In our example, we select **external**.
4. From the **Zone** list, select the appropriate zone. In our example, we select **example.com**
5. In the **Name** box, type a name for the Resource Record. Make sure the domain for which you are creating an MX record is shown, and note that it must end with a period.
6. In the **TTL** box, type a number of seconds. In our example, we type 500 (which is the default TTL for our zone).
7. From the **Type** list, select **MX**.
8. In the **Preference** box, type **10**. Preference is a numeric value for the preference of this mail exchange host relevant to all other mail exchange hosts for the domain. Lower numbers indicate a higher preference, or priority.
In a traditional DNS configuration, you would create multiple MX records with different priorities; however, since you're using GTM to provide true wide-area load balancing, it is only necessary to create a single record in this case.
9. In the **Mail Server**, enter the name of the Wide IP that you created in *Creating a wide IP on the GTM*. Make sure that this name also ends with a period. In our example, we type mail.example.com.
10. Click the **Finished** button (see Figure 3.7).

The screenshot shows a dialog box titled "Global Traffic » Records » New Resource Record...". Below the title bar is a section labeled "Record Configuration" with the following fields:

| | |
|-------------|-------------------|
| View Name | external |
| Zone Name | example.com. |
| Name | example.com. |
| TTL | 500 |
| Type | MX |
| Preference | 10 |
| Mail Server | mail.example.com. |

At the bottom of the dialog are three buttons: "Cancel", "Repeat", and "Finished".

Figure 3.7 *Creating a new Resource Record using ZoneRunner*

This concludes the BIG-IP GTM configuration. For more information on the BIG-IP GTM, see the GTM documentation.



4

Deploying the WANJet device with VMware ESX Servers

Configuring the WANJet device with VMware ESX Servers

This chapter describes how to configure WANJet devices to accelerate the transfer of files and reduce bandwidth utilization for ESX deployments. VMware allows administrators to easily copy virtual machine disk files (VMDKs) over a wide-area network (WAN) to duplicate an application or even an entire data center to a remote location. However, because VMDKs tend to be large -- typically several gigabytes up to many tens of gigabytes -- the time and bandwidth required for these file transfers are often excessive. By deploying F5 WANJets on either end of the WAN, you can accelerate the transfer of files and reduce overall bandwidth utilization, sometimes by dramatic amounts. Although compressibility and cacheability of individual VMDKs will vary greatly depending on operating system and other content and layout, it is not unusual to see a 3x or better improvement in transfer speeds while at the same time seeing a commensurate decrease in overall bandwidth use.

The F5 WANJet is a powerful, appliance-based solution that delivers fast, steady, predictable performance for all users and applications across a wide-area network (WAN). WANJet overcomes the limitations of wide area networks, helping with file transfers and data replication between data centers and branch offices. Through TCP optimizations, data compression, and intelligent byte-caching, WANJet greatly improves traffic performance that would otherwise suffer from excessive latency, constrained bandwidth, or unacceptable packet loss rates. Beginning in WANJet version 5.0, administrators have the option of using disk-based caching for added capacity, which is crucial for effective caching of large data sets such as VMDKs.

For more information on the WANJet device, see:
<http://www.f5.com/products/WANJet/>

Common Scenarios

VMware Virtual Infrastructure supports the use of virtual disks hosted via either a SAN (Storage Area Network) using FibreChannel or iSCSI, or via NAS (network-attached storage) using NFS. With those options, administrators have several options for replicating all of part of a set of virtual disks to remote destination.

- ◆ An NFS storage host might use NFS to duplicate files across the WAN to another remote storage host.
- ◆ A storage host that simultaneously supports both NFS and SMB/CIFS might use SMB/CIFS to duplicate files across the WAN to another remote storage host.
- ◆ A storage host might use FTP to duplicate files across WAN to another remote storage host.

- ◆ A SAN host might use implementation-specific native technologies (e.g. EMC SRDF, or Network Appliance SnapMirror) to duplicate entire volumes or volume changes to remote data centers.

Prerequisites and configuration notes

The following are prerequisites for this deployment

- ◆ A minimum of two WANJets (one at each end of the WAN); a redundant pair at either end is recommended.
- ◆ Each WANJet must be running version 5.0 or higher; we recommend at least 5.0.1. All WANJets must be on identical versions.
- ◆ Each WANJet must have at least one hard drive in order to enable disk-based caching.

Network Topology

WANjet can operate in either of two modes: In-Line or One-Arm.

- ◆ If configured as In-Line, WANjet acts as a transparent bridge. No routing changes are required to the storage hosts or other devices; however, the WANjet must be placed in the network in such a way that all WAN-destined traffic passes through it.
- ◆ If configured as One-Arm, WANjet acts as a router. In this case, devices may have to be re-configured to use the WANjet as a gateway to the remote network.

Pre-deployment tasks

In order to have a successful deployment, the replication system must already be functioning correctly, and meeting certain performance targets. We strongly recommend that the following checks be made prior to deployment.

Network traffic

- ◆ Document all applications sending traffic over the WAN Link and the protocol and port used. Common protocols for VMDK transfers are listed in the following table; it is possible that your configuration may be customized and different as a result.

| Protocol | Common Default Port(s) |
|-------------------|------------------------|
| FTP | 20, 21 (TCP) |
| NFS | 2049 (TCP) |
| SMB/CIFS | 139, 445 (TCP) |
| EMC SRDF | 1748 (TCP) |
| NetApp SnapMirror | 10566 (TCP) |

Table 4.1 Protocols and their common default ports

- ◆ Note the average and peak number of TCP connections for each of the applications sending traffic over the WAN link.
- ◆ Decide whether VMDK replication will be the only traffic to traverse the WANjet devices, or if all WAN traffic will do so. ***We highly recommend that all traffic between two remote locations is directed over the WANjet devices, so that they are aware of all connections and can optimize traffic more effectively as a result.***

Network configuration

Although WANJet is transparent to networking devices (such as routers and switches), these devices can impede the performance of the WANJet device if not configured correctly. For example, if a device is encrypting or compressing the traffic before it gets to WANJet, it will result in lower compression and data reduction performance of WANJet.

Document the following network components that sit between the storage hosts and the WAN routers on both the source and target sides.

- Routers
- Switches
- Hubs
- Encryption devices

Traffic that you intend to have the WANjet optimize should not be encrypted; for instance, storage hosts should not be using IPsec. WANjet can perform encryption on traffic it is optimizing, if desired.
- Compression devices

Traffic that you intend to have the WANjet optimize should not be previously compressed.

WAN Link

For optimal functionality, the WANJet device must know the actual available bandwidth, latency and packet loss averages over the WAN. Users often assume their WAN link meets certain parameters, when it may not. It is crucial to test all of these parameters under heavy traffic load and with the correct network priorities so that the measurements reflect the circuit that actual traffic will traverse. WAN links that seem to have certain characteristics under light load may have very different characteristics under heavy load. When the WANJet is accelerating traffic, it will often utilize all available bandwidth, so testing under load is critical.

Take note of the following network parameters while under maximum (or very heavy) load:

- Observed average and peak throughput in Megabits per second (Mbps).
- Observed average and peak latency via roundtrip time in milliseconds (RTT ms).
- Observed average and peak packet loss percentages.
- If possible, observed and average delay jitter is also a useful metric.

Configuring the WANJet devices for ESX devices

In this section, we configure the WANJet devices to accelerate the transfer of files and reduce bandwidth utilization for ESX deployments. Complete all of the following procedures.

Configuring the Operational Mode

On each WANjet device in this deployment, you must set up the Operational Mode.

To set the Operational Mode

1. In the navigation pane, expand **WAN Optimization** and then click **Operational Mode**. The Operational Mode screen opens.
2. Until all other WANJet configuration in the following sections is complete, leave the **Mode** set to **Inactive**.
3. In the **TDR-2 Storage Mode** section, click **Disk-Based Storage**.
4. In the **Topology** section, click the mode appropriate for your configuration. In our example, we click **In-Line**.
5. Click **Save** to save your changes.

| WAN Optimization >> Operational Mode | |
|--------------------------------------|----------------------------------|
| Operational Mode | |
| Mode | |
| Inactive: | <input checked="" type="radio"/> |
| Active: | <input type="radio"/> |
| TDR-2 Storage Mode | |
| Disk Based Storage: | <input checked="" type="radio"/> |
| Memory Based Storage: | <input type="radio"/> |
| Topology | |
| In-Line: | <input checked="" type="radio"/> |
| One-Arm: | <input type="radio"/> |
| Save Cancel | |

Figure 4.1 Configuring the Operational Mode

Configuring the Optimization Policy

The next step is to configure the Optimization Policy on each WANJet.

If you want to optimize all traffic across the WANJet devices, or are unsure of specific ports that your replication and other traffic may use, complete the following procedure. If you want to only optimize replication traffic, follow the subsequent procedure.

To optimize all traffic across the WANJet devices

1. In the navigation pane, expand **WAN Optimization**, and then click **Optimization Policy**.
2. In the Common section, in the Service Name column, click **All ports** (this link may say **All other ports** if you have already configured specific ports in this section).
The Edit Port/Service name box opens.
3. In the Edit Port/Service name box, from the **Processing Mode** list, select **Optimized**.
4. Make sure the **TDR-1**, **TDR-2**, and **Connection Intercept** boxes are all checked.
5. You can optionally check the **Encryption** box if you want WAN traffic to be encrypted.
6. Click the **OK** button.
The dialog box closes, and you return to the Optimization Policy page.
7. Click the **Save** button to commit your changes.

| Edit Port/Service Name | |
|------------------------|-------------------------------------|
| TCP - All other ports | |
| Processing Mode: | Optimized |
| TDR-1: | <input checked="" type="checkbox"/> |
| TDR-2: | <input checked="" type="checkbox"/> |
| Encryption: | <input type="checkbox"/> |
| Connection Intercept: | <input checked="" type="checkbox"/> |
| OK Cancel | |

Figure 4.2 Configuring the Optimization Policy for all traffic

If you wish to optimize only your replication traffic and each additional traffic flow for which you know details, use the following procedure.

To optimize only replication traffic

1. In the navigation pane, expand **WAN Optimization**, and then click **Optimization Policy**.

-
2. In the Common section, click the **Add** button.
 3. From the **Service Name** list, select the service, or type the port number in the first text field in the From Port section. (If there was a range of ports, you would enter the lowest number port in the first field, and the highest number port in the second.) In our example, we've selected NFS, which automatically populates **21** into the port field.
 4. From the **Processing Mode** list, select **Optimized**.
 5. Make sure the **TDR-1**, **TDR-2**, and **Connection Intercept** boxes are all checked.
 6. You can optionally check the **Encryption** box if you want WAN traffic to be encrypted.
 7. Click the **OK** button.
The dialog box closes, and you return to the Optimization Policy page.
 8. Repeat this procedure for each additional protocol, and then click **Save** to commit your changes.

Configuring the Tuning options

It is critical that the Tuning parameters be set correctly, or else traffic will not be optimized effectively.

To configure the Tuning options

1. In the navigation pane, expand **WAN Optimization**, and then click **Tuning**.
2. In the Bandwidth box, type the value in Bandwidth that accurately reflects your real WAN bandwidth, as noted the in *WAN Link*, on page 4. In our example, we typed **45** and selected **mb/s**, which is reflective of a dedicated T3 circuit.
3. In the **RTT** box, type the observed round-trip time as noted in *WAN Link*, on page 4. In our example, we type **80** milliseconds; your value will be specific to your circuit and the distance to your remote data center.
4. In most cases, make sure **Congestion Control** is not selected. If the WANJet congestion control is left checked (on), it interferes with the internal flow control mechanism used within some protocols, such as EMC Symmetrix SRDF. Typically, Congestion Control would be left On in situations where many simultaneous TCP connections were running through the WANJet (for example, over 1,000), which you might see in a branch office scenario. Since this is uncommon with data replication scenarios, and due to conflicts with some proprietary flow control mechanisms, we strongly recommend Congestion Control is Off.

5. Click **Save** to commit your changes.

| WAN Optimization >> Tuning | | |
|----------------------------|--------------------------|-------------|
| WANJet Tuning | | |
| Bandwidth: | 45 | mb/s |
| RTT: | 80 | msec |
| Congestion Control: | <input type="checkbox"/> | |
| | | Save Cancel |

Figure 4.3 Configuring the Tuning options

Modifying the Application QoS settings

The next step is to modify the Application QoS (Quality of Service) options. Application QoS bandwidth must be set to the actual bandwidth available to the WANJet devices, whether that is the entire link or only the Permanent Virtual Circuit (PVC) portion available. There is no need to reduce this setting to account for any overhead in the TCP protocol.

To modify the Application QoS settings

1. In the navigation pane, expand **WAN Optimization**, and then click **Application QoS**. The Application QoS screen opens.
2. From the Application QoS table, click the IP address of the remote WANJet appliance to which you want to apply an Application QoS policy. The Manage the Application QoS Settings of a Remote WANJet box opens.
3. In the **Link Bandwidth** box, type the bandwidth size of the link between the local WANJet and the remote WANJet. From the adjacent list, select **kb/s** or **mb/s**.
In our example, we are using an T3 link, so we type **45** and choose **Mb/s** (see Figure 4.4).

Note: If you have other application traffic sharing the same WAN link with your VMDK replication, it is possible to prioritize the Virtual Infrastructure traffic over the other traffic. An example would be prioritizing critical FTP-based VMDK replication traffic over less critical remote tape backup traffic. Refer to the WANJet User Guide for details on how to do this.

4. Click the **OK** button. The Manage the Application QoS Settings of a Remote WANJet box closes.
5. Click the **Save** button on the main Application QoS page.

Manage the Application QoS Settings of a Remote WANJet

Node Type:

WANJet IP:

WANJet Alias:

Link Bandwidth:

| Supported Subnet | Netmask | Alias | Status |
|------------------|---------------|-------|---------|
| 10.133.150.0 | 255.255.255.0 | | Enabled |

| Protocol | Service Name | Processing Mode | TDR-1 | TDR-2 | Encryption | Connection Intercept |
|----------|----------------------|-----------------|-------|-------|------------|----------------------|
| TCP | 20 (Active FTP data) | Optimized | Y | Y | N | Y |
| TCP | 21 (Ftp) | Optimized | Y | Y | N | Y |
| TCP | 445 (Microsoft-ds) | Optimized | Y | Y | N | Y |
| TCP | All other ports | Optimized | Y | Y | N | Y |

| Application QoS Policy | Bandwidth | Maximum |
|------------------------|-----------|---------|
| Default | 100% | 100% |

Figure 4.4 Configuring the Application QoS settings

Finalizing the Deployment

When you have completed all the preceding configuration procedures on the Local WANjet, repeat the steps for the Remote WANjet.

Activating the Optimization

When both WANjet devices are configured correctly, you must change the Operational Mode to Active on each WANjet.

To Change the Operational Mode

1. In the navigation pane, expand **WAN Optimization** and then click **Operational Mode**. The Operational Mode screen opens.
2. In the Mode section, click the **Active** button.
3. Click **Save** to save your changes.
After a short delay, the WANJet Links indicator at the top of the screen will turn green, indicating successful activation of the WANJet-to-WANjet link.

Additional Recommendations

The following sections contain additional considerations for deploying WANJet for VMDK replication.

Changes in the WAN environment

Anytime the WAN circuit available to the WANJets changes, such as the bandwidth or latency, all settings under the Application QoS and Optimization Policy must be adjusted correspondingly. Please note that latency input must correspond to the latency of the network under load. This means measuring latency while replication and other application traffic is flowing and tuning the WANJet accordingly.

Performance Charts

Leaving the real time performance charts portion of the WANJet GUI visible during while WANJet is actively operating can slightly reduce the overall performance of WANJet. For maximum performance, do not leave this tab open and in view for long periods of time.

SMB/CIFS Considerations

If SMB/CIFS is selected as the transport mechanism for VMDK transfers, a few additional steps must often be taken. Complete the following procedures.

Delayed Connection Acceptance

Use the following procedures to configure Delayed Connection Acceptance.

To configure Delayed Connection Acceptance

1. In the navigation pane, expand **WAN Optimization** and then click **Local WANjet**. The Local WANjet screen opens.
2. Select 'Settings for Delayed Congestion Acceptance'
3. Make sure ports **139** and **445** are entered, separated by a colon. (This is the default.)
4. If you make changes, click **Save** to apply the changes.

Disable SMB Signing

If your storage hosts are Windows-based (for instance, Windows 2003 Storage Server) and are using SMB/CIFS to transfer files between them, follow the instructions at <https://support.f5.com/kb/en-us/solutions/public/6000/200/sol6241.html?sr=539735> to turn off SMB signing on the storage hosts. Without disabling that setting, WANjet will not be able to optimize SMB/CIFS traffic.