



Deploying the BIG-IP v10.2 to Enable Long Distance VMotion with VMware vSphere

Table of Contents

Introducing the BIG-IP and VMware long-distance VMotion deployment guide	
Prerequisites and configuration notes	1-1
Product versions and revision history	1-2
Configuration details	1-2
Managing VM hosts, VM storage and client traffic between data centers during VMotion events	1-5
Components to be managed by automation	1-5
Configuration table	1-6
Configuring the primary data center BIG-IP system	1-9
Creating the VLANs	1-9
Creating self IP addresses	1-9
Configuring the application-specific objects on the BIG-IP system	1-10
Configuring the WAN optimization module	1-11
Configuring the Remote Endpoints	1-12
Confirming outbound connections are allowed	1-13
Advertising local networks	1-13
Configuring EtherIP	1-14
Configuring the VLAN group for failover BIG-IP pairs	1-15
Configuring the secondary data center BIG-IP system	1-17
Creating the VLANs	1-17
Creating the self IP addresses	1-17
Configuring the WAN optimization module	1-17
Configuring the remote endpoints	1-17
Confirming outbound connections are allowed	1-17
Advertising local networks	1-18
Configuring EtherIP	1-18
Configuring the BIG-IP GTM	1-19
Creating the data centers	1-19
Creating the monitor	1-19
Creating the GTM servers	1-19
Creating the GTM pool	1-20
Creating a wide IP on the GTM	1-21
Configuring the VMware infrastructure	1-22
Modifying the VMware ESX configuration	1-22
Appendix A: Test results	1-25
Testing methodology	1-25
Appendix B: Frequently asked questions and deployment considerations	1-27
Appendix C: Configuration worksheet	1-30

Introducing the BIG-IP and VMware long-distance VMotion deployment guide

Welcome to the BIG-IP system deployment guide for VMware vSphere™ VMotion™. This guide provides step by step instructions for configuring the BIG-IP system v10.2 with the WAN Optimization Module (WOM) for VMware long-distance VMotion implementations.

With this implementation, long-distance VMotion becomes possible between two geographically disparate data centers while the virtual machines being moved are active. Through the use of BIG-IP Global Traffic Manager (GTM) and the BIG-IP system's implementation of EtherIP, established user connections can also follow without interruption to the new datacenter.

New in the 10.2 release of BIG-IP WOM is the ability to fully use symmetric de-duplication, for even faster VMotion of storage and memory contents.

VMware VMotion technology (deployed in production by 70% of VMware customers according to a VMware customer survey from October 2008), leverages the complete virtualization of servers, storage and networking to move an entire running virtual machine with no downtime from one ESX server to another.

For more information on VMware vSphere VMotion, see <http://www.vmware.com/products/vmotion/>

For more information on the F5 devices included in this guide, see <http://www.f5.com/products/>.

You can also visit the VMware page of F5's online developer community, DevCentral, for VMware forums, solutions, blogs and more:

<http://devcentral.f5.com/Default.aspx?tabid=53&view=topics&forumid=46>.

To see test results of this deployment guide configuration, see *Appendix A: Test results*, on page 25.

To provide feedback on this deployment guide or other F5 solution documents, contact us at solutionsfeedback@f5.com.

Prerequisites and configuration notes

The following are general prerequisites for this deployment.

- ◆ You must have the WAN Optimization module licensed and provisioned on your BIG-IP systems, and be running version 10.2 or later.
- ◆ This guide includes configuration for the BIG-IP GTM. If you want to take advantage of the benefits of the BIG-IP GTM, you must have the GTM module licensed and provisioned on the BIG-IP system.

- ◆ Virtual IP Addresses in DNS controlled by Global Traffic Manager should have their time-to-live (TTL) records set to a minimum number of seconds in order to facilitate data center failover. Our recommendation is 15 seconds, but the times will vary depending on individual DNS architectures.
- ◆ There must be two BIG-IP systems running the WAN Optimization module, one as the local endpoint (primary) and one as the remote endpoint (secondary).
- ◆ Must have ESX VMotion and Storage VMotion licenses.
- ◆ VMware VMotion uses TCP port 8000. This port must be allowed to be initiated between the two data centers and between the two ESX servers.
- ◆ BIG-IP iSessions use TCP port 443. This port must be allowed to be initiated between the two data centers.
- ◆ VMware VMotion preserves all network settings of a virtual machine. While moving a machine between two data centers, the IP and network configuration for the migrated hosts between the two data centers must be identical. However, the infrastructure does not need to be part of the same layer 2 broadcast domain.
- ◆ See *Appendix B: Frequently asked questions and deployment considerations*, on page 27 for more information.

Product versions and revision history

Product and versions tested for this deployment guide:

Product Tested	Version Tested
BIG-IP system with the WAN Optimization module	v10.2
VMware vSphere VMotion	v4

Document Version	Description
1.0	New deployment guide

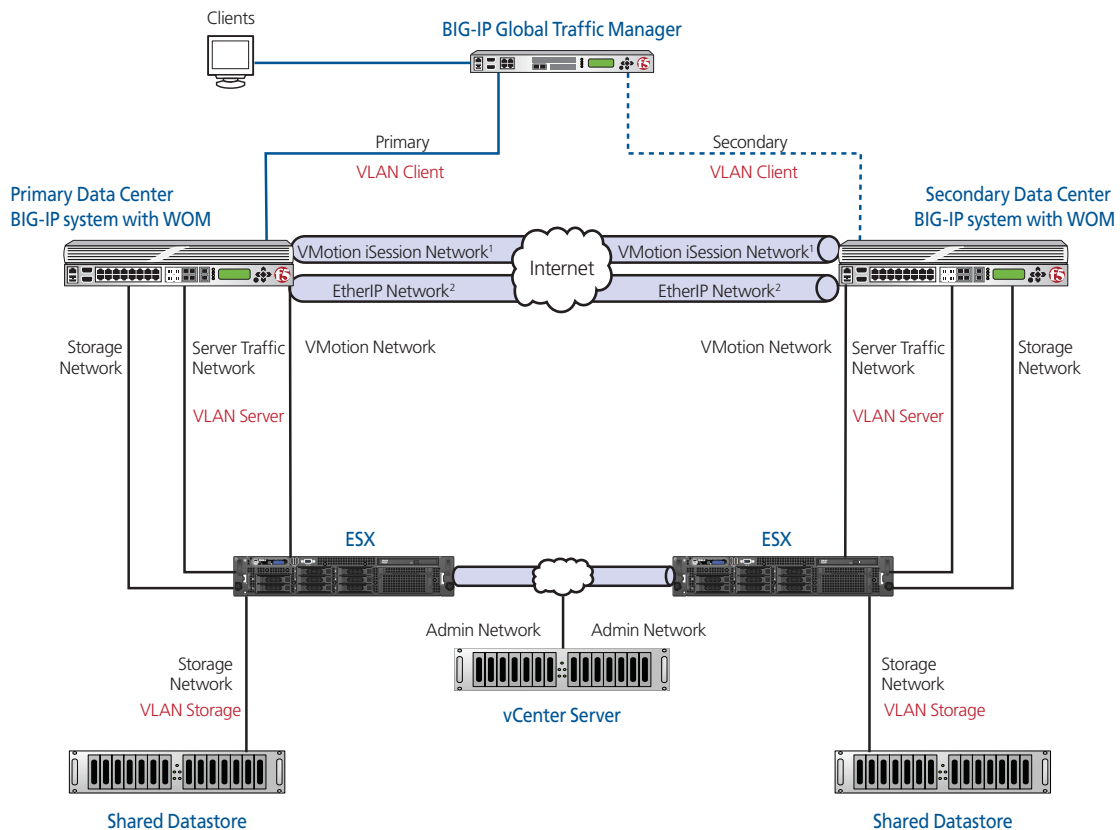
Configuration details

Figure 1, on page 3 is a logical configuration diagram highlighting the major components of a long-distance VMotion deployment. In our example, we have two data centers (Primary and Secondary) connected via a WAN (the actual WAN connectivity technology is not relevant in this architecture and may be IPSec, MPLS, Point-to-Point Frame Relay, or another technology).

Each data center contains nearly identical ESX and BIG-IP infrastructures. At a minimum, this configuration requires the ESX server must be VMotion compatible and the both BIG-IPs must support iSessions within the WAN Optimizations Module (WOM). Additionally, the VLANs, port groups and other immediate IP infrastructure with connected to the virtual machine must exist on each participating host. Upstream in both data centers, router and firewall configurations must also be consistent to allow client traffic.

In each data center, the ESX servers are configured as recommended by VMware, with client traffic residing on one dedicated physical network interface, the storage network on another interface, the administrative traffic on its own interface, and finally, the VMotion network on its own interface. ***By configuring VMWare ESX in this recommended manner, the VMotion network can have a separate TCP gateway address and therefore participate in long-distance encrypted and accelerated VMotion.***

An iSession tunnel is established between the BIG-IP systems in each data center. No further changes are required on other routers and switches to build this tunnel. The iSession tunnel uses TCP port 443, therefore that traffic has to be allowed through the firewall.



¹ The iSession network is encrypted and accelerated
² The EtherIP network is **not** encrypted or accelerated

Figure 1 Logical configuration example

Network Design

For long-distance VMotion traffic to succeed, on each ESX server we change the default gateway of the VMkernel managing VMotion to point to the BIG-IP system. This change is made through the Networking settings of the vCenter client. Optimization policies on the BIG-IP then recognize VMotion traffic and optimize it through iSession technology.

As in typical single data center deployments, vCenter Server is also used to manage Virtual Machines and VMotion in multiple data center deployments. However, unlike single data center deployments, the vCenter Server must be deployed to support two data centers in a long-distance VMotion deployment. There are two different modes you can use to accomplish this long-distance deployment.

In the first deployment mode, the vCenter Server can control servers in both data centers. In the second deployment example (not pictured in Figure 1) each data center has its own vCenter Server. After long-distance VMotion Migration, a particular host must be de-registered from the primary data center and re-registered in the secondary data center.

Finally, the last consideration in configuring long-distance VMotion is the strategy for managing Virtual Hosts. We have two recommendations in this regard. The first is to configure every server in your ESX cluster to participate in long-distance VMotion. If Dynamic Resource Scheduler (DRS) is turned on, this may not be possible.

The second recommendation is to use two dedicated migration hosts, one in each data center. These dedicated migration servers are responsible for long-distance VMotion and are the only ESX servers to have their configurations adjusted to participate in long-distance VMotion.

For this method to be successful, the migration hosts are setup to not participate in DRS (in their own cluster). To move a server from the primary data center to the secondary, first the virtual machine is moved to this ESX migration host, and then it is migrated over to the secondary data center. By using this procedure, dynamic VMotion, resource groups, and other provisioning configuration within existing ESX clusters do not need to be modified.

Managing VM hosts, VM storage and client traffic between data centers during VMotion events

This section contains information about the management of VM hosts (memory and storage) and client traffic between data centers while the VMotion events are occurring.

Components to be managed by automation

There are three components that can be scripted to more effectively manage a long distance VMotion.

- *Migrating Storage*
- *Migrating the virtual machine*
- *Using ratios to switch data center traffic with GTM*

Migrating Storage

For this solution, we recommend customers address the Storage VMotion first. For example, the movement of a 5 Gigabyte disk takes the longest amount of time (many minutes or hours) to complete. After this storage migration is completed, the memory portion of VMotion could be completed in a matter of seconds. This order of operation results in the least amount of time when disk I/O operations would be subjected to higher latency. Addressing storage needs has more than one possible solution:

- Using shared storage between two data centers with replication technologies from a third-party vendor.
- Using Storage VMotion and cross-mounting the shared storage between ESX hosts in both data centers.

Migrating the virtual machine

In order to manage hosts during VMotion events, the use of scripting and orchestration is recommended. The basic components for orchestration that can be used without additional expenditure are listed below. VMware also provides VMware Orchestrator, part of the vCenter Server Suite, which may be licensed for advanced automation.

- VMware's vSphere Web Services API
<http://www.vmware.com/support/developer/vc-sdk/>
- F5 BIG-IP's iControl API <http://devcentral.f5.com/Default.aspx?tabid=76>

Using ratios to switch data center traffic with GTM

We recommend the use of F5 iControl to dynamically manage the ratio or the cutover point for global traffic. This ensures traffic destined for one data center does not overwhelm an increasingly smaller number of hosts. To illustrate this, the following sections examine some typical long distance VMotion scenarios.

◆ Note

Because the final implementation depends on the automation or orchestration solution used by your implementation, we do not provide detailed procedures for configuring ratios. See the product documentation or DevCentral for more information.

Migrating a group of VMotion servers (2 or more)

For the migration of a group of hosts, management through scripting or orchestration, plus the use of the migration iRule and Priority Pool Activation is recommended to minimize client traffic disruption. As an example, if 10 hosts are to be evacuated from the primary data center to the secondary data center, the chain of events would be as follows:

- Administrative or automated decision is made to move the pool,
- Scripting or orchestration initiates the migration of storage from the primary data center to the secondary data center for the first host.
- Once storage is completed, the memory portion of the host is moved to the secondary data center.
- This process is repeated until 50% of the hosts are migrated at which point, GTM is instructed to direct traffic to the secondary data center,
- Host migration is completed, at which point, any traffic still arriving at the primary data center because of DNS or browser cache are retransmitted to the secondary data center through the use of Priority Pool Activation.

Configuration table

Because this implementation is relatively complex, the following configuration table shows the different VLAN, self IP and route settings in our example configuration.

Appendix C: Configuration worksheet, on page 30 contains a blank worksheet that you can fill out before beginning the BIG-IP configuration to streamline your deployment process.

Network	Primary Data Center	Secondary Data Center	Notes
iSession-WAN			
VLAN	vlan-isession-WAN1	vlan-isession-WAN2	
Self IP	10.133.57.141	10.133.58.141	Port Lockdown set to Allow None
Route	Static Route	Static Route	
iSession-LAN			
VLAN	vlan-isession-LAN1	vlan-isession-LAN2	
Self IP	10.133.59.245	10.133.60.245	Port Lockdown set to Allow Default
Route			
EtherIP			
VLAN	vlan-eip1	vlan-eip2	
Self IP	10.133.64.245	10.133.65.245	
Route	Static Route	Static Route	
Server			
VLAN	vlan-server1	vlan-server2	
Self IP	10.133.63.245	10.133.63.246	
Route (optional)			
Client			
VLAN	vlan-client1	vlan-client2	
Self IP	10.133.39.235	10.133.56.235	
Route (optional)			

The following is a description of the networks in the table above.

◆ **iSession-WAN**

This is the network that enables BIG-IP iSessions for deduplication, encryption and compression. The iSession network transfers Storage VMotion as well as Active State (memory) VMotion. The iSession network runs only between the two BIG-IP devices in the primary and secondary data center and needs to be routed properly to reach the destination each way.

◆ **iSession-LAN**

This is the network on the LAN side that terminates the VMotion traffic on either side. This Self IP will be the default gateway of VMware VMotion VMKernel.

◆ **EtherIP**

This is the network that enables connections to stay alive between both data centers. This tunnel does not provide encryption, compression or deduplication. The EtherIP network runs only between the two BIG-IPs in the primary and secondary data center and needs to be routed properly to reach the destination each way.

◆ **Server**

This network is where the VM servers are hosted by ESX. Notice the network has to be the same in both primary and secondary data center in order to pass VMware validation and to work with EtherIP (see FAQ section for additional details of VMware networking during VMotion).

◆ **Client**

This is the incoming client request traffic network. In this diagram, we are using private address space because there is upstream Network Address Translation (NAT) converting public IPs to our private space. In your scenario your client traffic network may be publicly routable IP space.

◆ **Tip**

Appendix C: Configuration worksheet, on page 1-30 contains a blank worksheet that you can fill out before beginning the BIG-IP configuration to streamline your deployment process.

Configuring the primary data center BIG-IP system

In this section, we configure the BIG-IP system located in the primary data center, including the WAN Optimization Module. After completing the primary data center BIG-IP system, we continue with *Configuring the secondary data center BIG-IP system*, on page 17.

◆ Note

*In this document, we typically refer to the **primary** and **secondary** data centers or BIG-IP systems. The WAN optimization module uses **local** and **remote**.*

Creating the VLANs

The first task is creating the VLANs on the BIG-IP system. For this configuration, you need to create the five VLANs outlined in the configuration table on page 8.

To create a VLAN

1. On the Main tab, expand **Network**, and then click **VLANs**.
The VLANs screen opens.
2. Click the **Create** button.
The new VLAN screen opens.
3. In the **Name** box, type a unique name for the VLAN. In our example for the iSession WAN VLAN we use **vmotion-issession-WAN1**.
4. In the **Tag** box, you can optionally type a tag. In our example, we leave this blank, and the BIG-IP LTM automatically assigns a tag.
5. In the Resources section, from the **Available** list, select the interface that will have access to tagged traffic, and add it to the **Untagged** box by clicking the Add (<<) button.
In our example, we select **1.14**.
6. Click the **Repeat** button.
7. Repeat steps 3-5 to create the other four VLANs, and then click the **Finished** button.

Creating self IP addresses

Self IP addresses are the IP addresses owned by the BIG-IP LTM system that you use to access the VLANs. The next task in this configuration is to create the five self IP addresses outlined in the configuration table above.

To create the self IP addresses

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 a static IP address that resides on the VLAN you created in the preceding procedure.
4. In the **Netmask** box, type the corresponding subnet mask.
5. From the **VLAN** list, select the appropriate VLAN you created in *Creating the VLANs*.
6. For the iSession network only: From the **Port Lockdown** list, select **Allow None**. In our example, this is **vlan-isession-WAN1**.
7. Click the **Repeat** button.
8. Repeat steps 3-6 for each self IP address, and then click **Finished**.

Configuring the application-specific objects on the BIG-IP system

Because this guide is not specific to any particular application, we recommend you use the deployment guide specific to your application. A list of deployment guides can be found at

<http://www.f5.com/solutions/resources/deployment-guides/>

However, there are some guidelines you must follow when configuring the application-specific pools and virtual servers on the BIG-IP when using this deployment guide to enable long-distance VMotion.

◆ **BIG-IP Pool members**

Pool members involved in Long Distance VMotion should be the same in both the primary and secondary data centers.

In our example, we have pool members in the 10.133.63.x/24 network (10.133.63.50, 10.133.63.51, and 10.133.63.52) in the primary data center. We use these exact same IP addresses when configuring the BIG-IP pool in the secondary data center.

◆ **BIG-IP virtual servers**

The BIG-IP virtual servers can be any public or private address space the infrastructure needs. Note that in our example, the server network uses 10.133.39.x and 10.133.56.x.

IMPORTANT: Enable **SNAT Automap** in order to ensure the pool member routes back to the Virtual Server from which it received traffic. The SNAT Automap setting is found on the BIG-IP virtual server configuration page.

For more information on SNAT, see the product documentation.

Configuring the WAN optimization module

In this section, we configure the WAN optimization module (WOM). The WAN optimization module allows you to encrypt and accelerate data between BIG-IP devices, accelerate applications across the WAN, and much more.

We recommend creating your own certificates for iSession communication. We recommend using certificates for iSession, as it gives you granular control over SSL certificates. For more information about SSL certificates on the BIG-IP system, see the product documentation available at <https://support.f5.com/> or the online help.

One of the options in configuring the WAN optimization module is the choice to use Dynamic Discovery. The benefit of dynamic discovery is that it reduces configuration complexity. However, when dynamic discovery is used, the BIG-IP currently disables iSession routing in order to prevent inadvertent routing loops. In our environment, dynamic discovery is allowed, but care was taken to ensure iSession routing was enabled.

To configure the WOM module

1. On the Main tab of the local BIG-IP system, expand **WAN Optimization**, and then click **Quick Start**. The Quick Start configuration screen opens.
2. In the **WAN Self IP Address** box, type the BIG-IP self IP address you provisioned the WAN Endpoint.
3. From the **Discovery** list, select **Enabled**.
4. In the **LAN VLANs** section, from the **Available** list, select the Server VLAN you created in *Creating the VLANs*, on page 9, and then click the Add (<<) button. In our example, we click **vlan-server1**.
5. In the **WAN VLANs** section, from the **Available** list, select the iSession VLAN you created in *Creating the VLANs*, on page 9, and then click the Add (<<) button. In our example, we click **vlan-isession-WAN1**.
6. In the Authentication and Encryption section, from the **Outbound iSession to WAN** list, do one of the following:
 - If you already have already created a Server SSL profile with certificate and key information specific for iSession, select it from the list and continue with step 7.
 - If you have not created a Server SSL profile specific to the iSession tunnel, click the Add (+) button to create one. The New Server SSL Profile page opens.
 - a) In the **Name** box, type a name for this profile. In our example, we type **iSession-VMotion-OutboundSSL**.
 - b) From the **Certificate** and **Key** lists, select the appropriate certificate and key.

- c) Click the **Finished** button. You return to the WOM configuration page.
 - d) Select the profile you just created from the list.
7. From the **Inbound iSession to WAN** list, do one of the following:
 - If you already have already created a Client SSL profile with certificate and key information, select it from the list and continue with step 8.
 - If you have not created a Client SSL profile, click the Add (+) button to create one. The New Client SSL Profile page opens.
 - a) In the **Name** box, type a name for this profile. In our example, we type **iSession-VMotion-InboundSSL**.
 - b) From the **Certificate** and **Key** lists, select the appropriate certificate and key.
 - c) Click the **Finished** button. You return to the WOM configuration page.
 - d) Select the profile you just created from the list.
8. From the **Application Data Encryption** list, we strongly recommend selecting **Enabled** from the list. VMware does not encrypt VMotion data.
9. In the **Create Optimized Applications** section, check the box for **VMware VMotion**. You see a green checkmark next to VMware VMotion and Data Encryption should be set to Enabled.
10. Click the **Apply** button.

Configuring the Remote Endpoints

The next task is to configure the Remote Endpoints on the BIG-IP WAN optimization module.

To configure the remote endpoints

1. On the Main tab of the local BIG-IP system, expand **WAN Optimization**, and then click **Remote Endpoints**.
2. Click the **Create** button.
3. In the **IP Address** box, type the Self IP address for iSession in the secondary data center. In our example, we type **10.133.58.141**.
4. Leave all other settings at the defaults.

-
5. Click the **Finished** button. The remote endpoint is added to the list.

◆ **Note**

Ensure that your BIG-IP system has the appropriate route that indicates how traffic should reach the remote BIG-IP system. For example, if your system does not use a default gateway to reach the remote BIG-IP iSession endpoint, create a static route indicate the gateway to be used. See BIG-IP product documentation on how to create routes.

Confirming outbound connections are allowed

The next task is to confirm that outbound connections are allowed.

To confirm outbound connections are allowed

1. On the Main tab of the local BIG-IP system, expand **WAN Optimization**, and then click **Remote Endpoints**.
2. In the table, click the IP address for the remote Endpoint you just created.
3. In the Outbound iSession to WAN section, make sure there is a check in the **Outbound Connections** box. If there is not, check the box.
4. Click **Update**. You return to the Remote Endpoints list.
5. In the Remote Endpoints table, click a check in the box next to the IP address of the Remote Endpoint you modified, and then click the **Manual Save** button.

Advertising local networks

The next task is to advertise the local networks.

To advertise the local networks

1. On the Main tab of the local BIG-IP system, expand **WAN Optimization**, and then click **Advertised Routes**.
2. Click the **Create** button.
3. In the **Address** box, type the IP Address of the Client Network. In our example, this is the *vlan-client1* vlan, so we type **10.133.39.0**.
4. In the **Netmask** box, type the corresponding subnet netmask. In our example, we type **255.255.255.0**.
5. Click the **Finished** button.

Configuring EtherIP

The next task is to configure EtherIP on the BIG-IP system. The EtherIP configuration must be done using the command line. There are two procedures for configuring EtherIP; configuring the EtherIP tunnel, and configuring a VLAN group.

Creating the EtherIP tunnel

The first procedure in configuring EtherIP is to create the EtherIP tunnel on the BIG-IP system. This must be done using the command line.

To configure the EtherIP tunnel

1. On the BIG-IP system, start a console session.
2. Type a user name and password, and then press Enter.
3. At the system prompt, type **tmsh** and then press Enter. This opens the Traffic Management shell.
4. Type **net tunnel**, and press Enter.
5. Use the following syntax to create the tunnel:

```
create tunnel <tunnel name> profile etherip local-address <local_self_ip_address>  
remote-address <remote_self_ip_address>
```

The self IP addresses that you specify are those that you created for EtherIP VLAN (**vlan-eip1** and **vlan-eip2** in our example) on both the local and the remote BIG-IP system.

In our example, we type:

```
create tunnel eip profile etherip local-address 10.133.64.245 remote-address 10.133.65.245
```

6. Type **save / sys config**, and press Enter.
7. To exit the shell, type **quit**, and press Enter.

Creating the VLAN group

To complete the EtherIP tunnel configuration, you must create a VLAN group. This VLAN group associates the EtherIP traffic with the BIG-IP virtual server traffic. This allows the BIG-IP system to recognize where servers are located, either “locally” or “remotely” (via the EtherIP interface). To create the VLAN group, you must use the command line, however once it is created, it can be edited using the BIG-IP Configuration utility (GUI).

To create the VLAN group

1. On the BIG-IP system, start a console session.
2. Type a user name and password and then press Enter.
3. At the system prompt, type **tmsh** and then press Enter. This opens the Traffic Management shell.

4. Type **net vlan-group**, and then press Enter.

5. Use the following syntax to create the VLAN group:

```
create <vlangroup-name> mode <mode> members add { <vlan-name> } members add <vlan-name> }
```

In our example, we associate the **vlan-eip** and **vlan-server** with this VLAN group, so we type

```
create vg-eip mode transparent members add { vlan-eip1 } members add { vlan-server1}
```

6. Type **save / sys config**, and press Enter.

7. To exit the shell, type **quit**, and press Enter.

8. Open the BIG-IP Configuration utility using a web browser. In the following steps, we verify that the VLAN group was successfully added.

9. On the Main tab of the navigation pane, expand **Network** and then click **VLANs**.

10. From the menu bar, click **VLAN Groups**, and then click **List**.

11. Click the Name of the VLAN group you just created. Verify it has two members; the EtherIP tunnel and the VLAN for your servers.

Configuring the VLAN group for failover BIG-IP pairs

When setting up BIG-IP in a failover pair, the Media Access Control (MAC) Masquerade feature should be used to insure seamless traffic failover between the Active and Standby BIG-IP devices. In order to configure MAC Masquerade, you need to create a unique MAC address for this VLAN Group. When there is a failover event between the BIG-IP devices, while EtherIP is in use, traffic will not be disrupted.

You must select two unique MAC address to be shared between your active and standby BIG-IP devices (One per data center; the standby BIG-IP device uses its own MAC address when it is not active). Selecting a unique and locally administered Media Access Control (MAC) address is important to insure there are no overlaps with any other MAC address on your networks. F5 solution 3523

<https://support.f5.com/kb/en-us/solutions/public/3000/500/sol3523.html> describes this process. The convention to designate an address as locally administered, recommended by F5, is to flip the second to last bit of the first byte of the MAC address to one. The following table from SOL3523 illustrates how to do this. See the solution for more information.

Pre-assigned MAC address	First byte	Local bit	Flipped local bit	New first byte	Locally administered MAC address
00:01:D7:01:02:03	00	00000000	00000010	02	02:01:D7:01:02:03
01:01:D7:01:02:03	01	00000001	00000011	03	03:01:D7:01:02:03
08:01:D7:01:02:03	08	00001000	00001010	0A	0A:01:D7:01:02:03

Table 1 MAC address conversion table from Ask F5

To determine your own unique, locally administered MAC address

1. On the Main tab of the navigation pane, expand **Network** and then click **Interfaces**.
2. In the fourth column, note the physical MAC address associated with the physical interface used for EtherIP traffic. In our case it is: **0:1:d7:92:c0:c4**.

In this case, the first byte is 00 (listed as a single zero in the display).

The local bit in this case is 00000000 (eight zeros), we flip the second to last bit and we now have 00000010. Our new first byte is now 02.

3. In our example, we replace our new first byte and end up with: **02:01:d7:92:c0:c4**. We use this in step 5 of the following procedure.

This scheme guarantees the MAC address is always unique.

To configure MAC Masquerade

1. On the Main tab of the navigation pane, expand **Network** and then click **VLANs**.
2. From the menu bar, click **VLAN Groups**, and then click **List**.
3. Click the name of the VLAN group you created in *Creating the VLAN group*, on page 1-14.
4. Make sure the **Bridge in Standby** box is not checked.
5. In the **MAC Masquerade** box, type the unique MAC address you calculated in the preceding procedure.
6. Click **Update**.

Repeat this MAC Masquerade section for your secondary site, insuring that you create a brand new and unique locally administered MAC address using the instructions here.

This concludes the EtherIP configuration.

◆ Note

*If your configuration is static and dedicated to VMotion, we recommend you disable **Auto Discovery**. After you have completed all of the procedures above and verified that VMotion is traversing iSession, disable Auto Discovery by expanding **WAN Optimization**, clicking **Quick Start** and then, from the **Auto Discovery** list, selecting **Disable**. Click **Apply** when finished.*

In the next section, we configure the secondary data center BIG-IP system.

Configuring the secondary data center BIG-IP system

In this section, we configure the secondary data center (or remote) BIG-IP system. The procedures are nearly identical to those in the preceding section, we refer back to those sections instead of repeating the entire procedures.

Creating the VLANs

Follow the procedure *Creating the VLANs*, on page 9 to create five VLANs on the secondary data center BIG-IP system. Refer to the *Configuration table*, on page 6 for our examples.

Creating the self IP addresses

Follow the procedure *Creating self IP addresses*, on page 9 to create five self IP addresses on the secondary data center BIG-IP system. Refer to the *Configuration table*, on page 6 for our examples.

Configuring the WAN optimization module

Follow the procedure *Configuring the WAN optimization module*, on page 11 to configure the WAN optimization module in the secondary data center. Note the following minor changes:

- In step **4**: in our example, we select **vlan-server2**.
- In step **5**: in our example, we select **vlan-isession-WAN2**.
- In step **6a**: give the profile a unique name.
- In step **7a**: give the profile a unique name.

Configuring the remote endpoints

Follow the procedure *Configuring the Remote Endpoints*, on page 12 configure the remote endpoints on the secondary data center BIG-IP system with the following change:

- In step 3, type the Self IP address for iSession in the primary data center. In our example, we type **10.133.57.141**.

Confirming outbound connections are allowed

Follow the procedure *Confirming outbound connections are allowed*, on page 13 to confirm outbound connections are allowed on the secondary data center BIG-IP system.

Advertising local networks

Follow the procedure *Advertising local networks*, on page 13 to advertise the local networks on the secondary data center BIG-IP system with the following change:

- In step 3, in the **Address** box, type the IP Address of the Client Network. In our example, this is the *vlan-client2* vlan, so we type **10.133.56.0**.

Configuring EtherIP

Follow the procedure *Configuring EtherIP*, on page 14 to configure EtherIP on the secondary data center BIG-IP system with the following changes:

- ◆ For the *Creating the EtherIP tunnel*, on page 14:
 - In step 5, use the following syntax to create the tunnel (keeping in mind that the local address is now in the secondary data center):

```
create tunnel <tunnel name> profile etherip local-address <local_self_ip_address>  
remote-address <remote_self_ip_address>
```

The self IP addresses that you specify are those that you created for EtherIP VLAN (**vlan-eip1** and **vlan-eip2** in our example) on both the local and the remote BIG-IP system.

In our example, we type:

```
create tunnel eip profile etherip local-address 10.133.65.245 remote-address 10.133.64.245
```

- ◆ For the *Creating the VLAN group*, on page 14
 - In step 5, we use the following as our example:

```
create vg-eip mode transparent members add { vm-eip2 } members add { vlan-server2}
```

This completes the secondary data center BIG-IP configuration.

Configuring the BIG-IP GTM

F5's Global Traffic Manager must be configured to direct traffic to the correct LTM virtual server. In our example, we send all traffic to the local data center, unless the utilization alarms we configure are triggered.

Creating the data centers

In this task you need to create two data centers, called Local and Remote respectively, that correspond to your physical data centers.

To create the data centers

1. On the Main tab of the navigation pane, expand **Global Traffic** and click **Data Centers**. The main screen for data centers opens.
2. Click the **Create** button. The New Data Center screen opens.
3. In the **Name** box, type a name for this data center. In our example, we type **Local**.
4. Complete the rest of the configuration as applicable for your deployment.
5. Click the **Finished** button. Repeat this procedure for the Remote data center.

Creating the monitor

The next step is to create an HTTP monitor.

To create the 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 **VM-http-monitor**.
4. From the **Type** list, select **HTTP**.
5. Configure the options as applicable for your deployment.
6. Click the **Finished** button. The new monitor is added to the list.

Creating the GTM servers

The next task is to create servers on the BIG-IP GTM system.

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 **Local-BIG-IP**.
4. Configure the properties of the server as applicable for your deployment.
5. In the Health Monitors section, from the **Available** list, select the name of the monitor you created in *Creating the monitor*, on page 19, and click the Add (<<) button. In our example, we select **VM-http-monitor**.
6. Click the **Finished** button

Creating the GTM pool

The next task is to create a pool on the BIG-IP GTM system that includes the LTM virtual server in the local data center, and one that includes the LTM virtual server in the remote data center. The remote data center pool should be Disabled after creation.

To create a GTM pool

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 **Local_pool**.
4. In the Health Monitors section, from the **Available** list, select the name of the monitor you created in *Creating the monitor*, on page 19, and click the Add (<<) button. In our example, we select **VM-http-monitor**.
5. In the Load Balancing Method section, choose the load balancing methods from the lists appropriate for your configuration.
6. In the Member List section, from the **Virtual Server** list, select the virtual server you created for the application, and click the **Add** button. Note that you must select the virtual server by IP Address and port number combination. In our example, we select **10.133.39.51:80**.
7. Configure the other settings as applicable for your deployment
8. Click the **Finished** button.

Creating a wide IP on the GTM

The final step in the GTM configuration is to create a wide IP that includes both newly-created pools, and uses the fully qualified domain name (FQDN) you wish to use for the application. In our example, we use **vmhttp.siterequest.com**.

To create a wide IP

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 **vmhttp.siterequest.com**.
4. From the **State** list, ensure that **Enabled** is selected.
5. From the Pools section, from the Load Balancing Method list, select a load balancing method appropriate for your configuration.
6. In the Pool List section, from the **Pool** list, select the name of the pool you created in *Creating the GTM pool*, on page 20, and then click the **Add** button. In our example, we select **Local_pool**. Repeat this step for the remote pool.
7. All other settings are optional, configure as appropriate for your deployment.
8. Click the **Finished** button.

This completes the basic GTM configuration. For more advanced GTM configuration options, see the BIG-IP GTM documentation.

Configuring the VMware infrastructure

In this deployment guide, we assume you already have your VMware VMotion implementation up and running. However, there are some modifications you need to make to the VMware configuration for the configuration in this guide to work properly.

Modifying the VMware ESX configuration

The ESX servers should be configured to have a VMkernel Port for VMotion. This VMkernel port, on an ESX virtual switch should be bound to a unique physical adapter. Each ESX server should have a shared stored device mounted via iSCSI or NFS that both ESX servers can mount. Thus, for testing, storage does not become a gating factor.

Modifying the VMkernel default gateway

The next task is to modify the default gateway on the VMkernel for VMotion to the self IP address you created in *Creating self IP addresses*, on page 9.

To modify the VMkernel default gateway

1. Open the VMware vSphere client, and select the appropriate ESX server host.
2. Click the Configuration tab.
3. In the Hardware box, click **Networking**.
4. From the Networking list, locate the Virtual Switch that contains the VMotion kernel.
5. Click the **Properties** link.




Figure 2 Properties link of the Virtual Switch

6. Click to highlight **VMkernel** and then click the **Edit** button. The VMkernel properties box opens.
7. Click the IP Settings tab.
8. Click the **Edit** button next to **VMkernel Default Gateway**. The DNS and Routing Configuration box opens.
9. In the **Default Gateway** box, type the IP address of the self IP on the BIG-IP device.

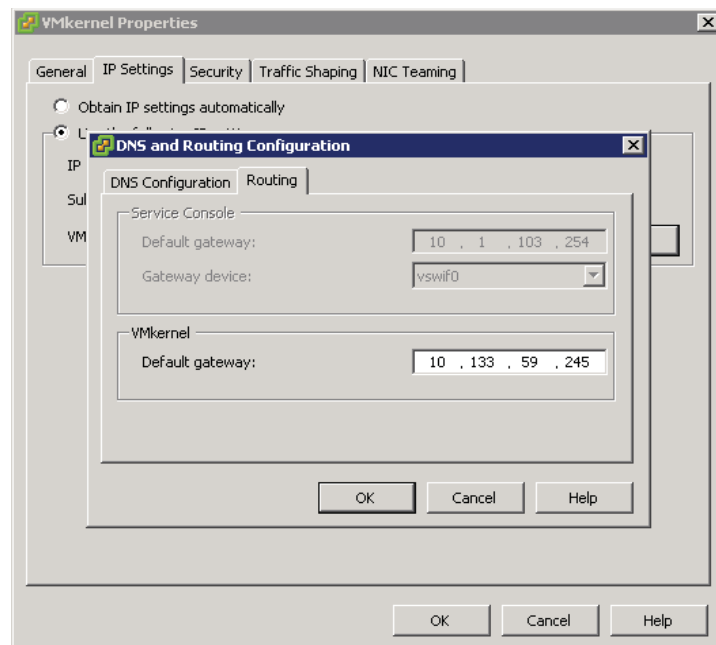


Figure 3 Modifying the default gateway

10. Click the **OK** button, and then click **OK** and **Close** to close all the open windows.

The same procedure must be performed on additional ESX servers in both data centers. The VMkernel default gateway in each location should be on a local network.

Binding the ESX devices to a specific vmknics

The final task is to bind the ESX machine to a specific vmknics.

To modify the VMkernel default gateway

1. Open the VMware vSphere client, and select the appropriate ESX host.
2. Click the Configuration tab.

3. In the Software box, click **Advanced Settings**. The Advanced Settings window opens.
4. From the left navigation tree, click **Migrate**.
5. In the **Migrate.BindToVmknics** row, type **1** in the box.
6. Click the **OK** button.

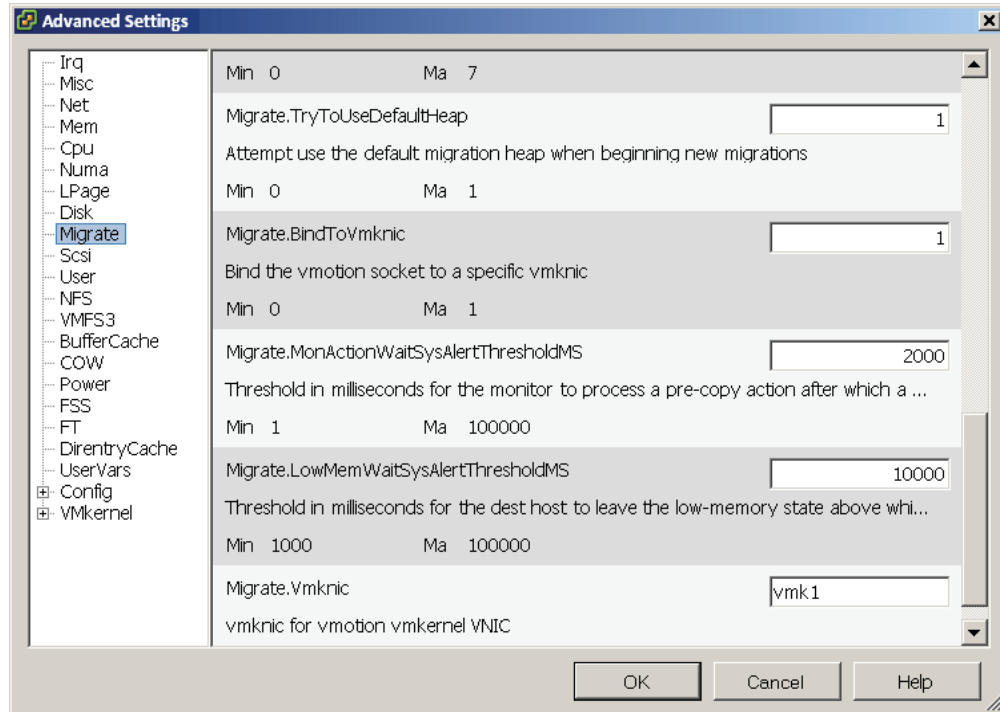


Figure 4 Modifying the *Migrate.BindToVmknics* option

This completes the deployment guide configuration.

Appendix A: Test results

Testing of long-distance VMotion was carried out in F5's technology center using ESX Server version 4.0, BIG-IP version 10.1 and virtual machines running both Windows XP and Linux. The virtual machines were configured as follows:

- ◆ 1 Gig of RAM
RAM fully consumed by a content management system; the machine was swapping to disk. The amount of active memory moved across the network was 1 gigabyte.
- ◆ 1 CPU
Fully utilized (0% idle, with some processes blocked).
- ◆ 10 Gigabytes of disk space
About 50% utilized. Note that the results quoted are primarily for RAM contents, but the same acceleration is seen on Storage VMotion.

It is important to note that the virtual machines were fully loaded (as described above) during the VMotion events. Active virtual machines take more time and resources to migrate than idle virtual machines.

Testing methodology

VMotion testing was conducted by initiating VMotion using VMware's vSphere Web Services API while a load test was running to fully use all resources. Part of the test methodology was to insure that there was minimal or no user disruption during the HTTP based test against the content management system.

The result using various network conditions follow. The first result demonstrates a large amount of bandwidth and low latency, the second result is with relatively large bandwidth but larger latency. It is evident that even a difference of 20 ms causes large slow-downs for un-aided VMotion. Finally, in the last scenario, there is a fair amount of bandwidth but much higher latency.

The decision on which Storage VMotion to use depends on the type of application, the allowable latencies for users and the distance between the two data centers.

The 622 (OC12) and 1000 results, shaded in the following table, are network conditions that VMware has also tested.

Network Conditions			BIG-IP with iSessions - average time in seconds	No acceleration - average time in seconds
<i>Mbps</i>	<i>RTT Latency</i>	<i>Packet Loss</i>		
45 (T3)	100 ms	0%	215 seconds	823 seconds
100	25 ms	0%	78 seconds	370 seconds
155 (OC3)	100 ms	0%	209 seconds	805 seconds
622 (OC12)	40 ms	0%	117 seconds	357 seconds
1000	20 ms	0%	38 seconds	132 seconds

Table 2 Typical migration times of an active virtual machine with 1 gigabyte of memory

Notes:

- For the Gigabit LAN case tests, Jumbo Frames were not turned on.
- Mbps is Megabits per second
- RTT = Round Trip Time

Appendix B: Frequently asked questions and deployment considerations

Q: Doesn't VMotion require ESX hosts to share a layer 2 bridge? How does this work over a traditional WAN, like the Internet?

A: A common misconception is that VMotion requires a layer 2 bridge to work across a WAN. However, the only technical requirement from a network standpoint for VMotion to succeed is that the network from the guest VM perspective remain identical. This means the guest IP address remains unchanged, and all port groups which touch the guest exist in the source and target ESX hosts. The VMotion traffic itself uses the VMkernel port of ESX, and this does not have to be identical on each host. It is through the VMkernel port and default gateway (which is not shared by the guest VM) that we route traffic across the iSession tunnel.

Q: Can you summarize what IP addresses are different, and why?

A: The following tables summarize the key network IP addresses.

IP address	Description	Different for each data center?
LTM virtual server	Public IP which is used for client connections between clients and the LTM. The GTM determines which LTM virtual server to direct clients to when they initialize new application connections.	Yes
VM guest IP	The IP address the guest uses to receive and respond to client requests.	No
ESX VMkernel IP	Used for VMotion and Storage VMotion traffic.	Yes
ESX VMkernel default gateway	The gateway used to route VMotion traffic between hosts.	Yes. Specifically, this value will be a self IP on the local LTM of each data center.

Q: The guide mentions use of a dedicated migration host in each data center to transition from one vCenter to another. Can you elaborate?

A: An alternative deployment scenario would leverage a dedicated host in each data center to be used for long-distance VMotion, analogous to a dedicated host in VMware Fault Tolerant deployments. This host would not be a member of any DRS or HA resource pools, nor would it use any Distributed Virtual Switches you may have in your cluster. In this scenario, the dedicated hosts in each data center would only have to be able talk to the other hosts on the Service Console network (or Management Network on ESXi), and have the same port groups which are accessed by the VM configured on the standard switch(es) of that host (as required by VMotion).

The work flow of migrating a virtual machine from, for example, a DRS cluster from one data center into a DRS cluster in the other data center would work as follows:

1. Both transition hosts would be initially managed by the vCenter in the primary data center.
2. Configure the hosts with access to a datastore in both the primary and secondary data centers.
3. VMotion the VM from the DRS resource pool to this dedicated host.
4. Storage VMotion, then VMotion the VM from the local dedicated host/datastore to the remote dedicated host/datastore.
5. De-register the secondary host from vCenter in the primary site.
6. Register this host with the vCenter in the secondary site.
7. VMotion the VM from this host into the local DRS resource pool in the target data center.

Q: Do any MAC addresses change during VMotion? What about ARPs and reverse ARPs?

A: No. A key principle of VMotion is that the networking stack from the guest perspective remain exactly the same before and after a VMotion. In a typical LAN scenario, a reverse ARP is issued at the completion of the VMotion in order to tell the local router that traffic bound for the VM has moved guests, and to direct that traffic to the new MAC address of the destination host. This is necessary because in a non-F5 enhanced VMotion event, both hosts are on the same broadcast domain.

However, in this solution, the guest has moved from a host in one physical data center to another. The task of routing traffic bound for the guest is managed not by a single local router, but by GTM and LTM. When the guest arrives in the secondary data center, inbound connections continue to reach the guest VM, because GTM and the LTMs are aware of the guest's location dynamically, and will route those inbound connections to the correct data center where the guest lives at that moment in time. MAC addresses do not change on the guest nor the hosts.

Q: What are the optimal use cases for F5's long-distance VMotion solution?

A: A key element of this solution is the transparent redirection of inbound connections between clients and application VMs. Migrating web applications between data centers is the ideal use case for this type of solution, as web applications have many short lived connections between clients and servers. Web applications are almost always session based, meaning once a user begins using a web application, it is important that all requests from that user persist to the same VM. Should that VM migrate from one data center to another, requests from an existing user session must continue to reach the same VM. The F5 solution meets these requirements transparently and effectively, making it an ideal use case.

Applications that have long-lived persistent connections, such as SSH, telnet, or streaming media, are not good use cases. Similarly, any applications that are highly transactional, such as database applications, are not good use cases for the solution. Attempting to perform a Storage VMotion (whether local or over long-distance) of a database is not recommended and such use cases are better addressed using database replication solutions from storage vendors, which are purpose built for moving and synchronizing highly transactional data between sites.

Q: What are some suggested strategies for automating this solution?

A: One of the key benefits of both VMware and F5 solutions is the ability to automate complex tasks through published APIs. Automation decreases the risk of human error, simplifies complexity, and reduces operating costs associated with a task by streamlining workflow.

Fortunately, this solution lends itself quite well to automation.

Many organizations already have workflow engines in their environment to automate tasks. Others develop scripts in-house for common functions. In either scenario, the discreet steps of executing a long-distance VMotion can be programmatically executed using the VMware vSphere Web Services API:

1. Execute the Storage vMotion
2. Execute the VMotion
3. (optionally) De-register host with vCenter in the primary data center
4. (optionally) Register host with vCenter in secondary data center.

For further discussion on VMware or VMotion, visit the VMware forums on DevCentral:

devcentral.f5.com/Default.aspx?tabid=53&view=topics&forumid=46

Appendix C: Configuration worksheet

The following is a blank worksheet you can use to assist you in the BIG-IP configuration. For our example, see *Configuration table*, on page 6.

Network	Primary Data Center	Secondary Data Center	Notes
iSession-WAN			
VLAN			
Self IP			Port Lockdown set to Allow None
Route			
iSession-LAN			
VLAN			
Self IP			Port Lockdown set to Allow Default
Route			
EtherIP			
VLAN			
Self IP			
Route			
Server			
VLAN			
Self IP			
Route (optional)			
Client			
VLAN			
Self IP			
Route (optional)			

The following is a description of the networks in the table above.

- ◆ **iSession-WAN**

This is the network that enables BIG-IP iSessions for deduplication, encryption and compression. The iSession network transfers Storage VMotion as well as Active State (memory) VMotion. The iSession network runs only between the two BIG-IP devices in the primary and secondary data center and needs to be routed properly to reach the destination each way.

- ◆ **iSession-LAN**

This is the network on the LAN side that terminates the VMotion traffic on either side. This Self IP will be the default gateway of VMware VMotion VMKernel.

- ◆ **EtherIP**

This is the network that enables connections to stay alive between both data centers. This tunnel does not provide encryption, compression or deduplication. The EtherIP network runs only between the two BIG-IPs in the primary and secondary data center and needs to be routed properly to reach the destination each way.

- ◆ **Server**

This network is where the VM servers are hosted by ESX. Notice the network has to be the same in both primary and secondary data center in order to pass VMware validation and to work with EtherIP (see FAQ section for additional details of VMware networking during VMotion).

- ◆ **Client**

This is the incoming client request traffic network. In this diagram, we are using private address space because there is upstream Network Address Translation (NAT) converting public IPs to our private space. In your scenario your client traffic network may be publicly routable IP space.