VMware vCenter Site Recovery Manager

Welcome to the BIG-IP deployment guide for VMware Site Recovery Manager (SRM). This guide provides procedures for configuring the BIG-IP Local Traffic Manager (LTM), Global Traffic Manager (GTM), and Access Policy Manager (APM) with VMware SRM. This document also presents guidance on how to configure the virtual infrastructure that will participate in site recovery with SRM.

VMware vCenter™ Site Recovery Manager extends VMware vCenter to enable automated cold migrations (i.e., a shut down state) of virtual machines (VMs) from one data center to another. Together, SRM and the F5 BIG-IP can simplify the implementation and mitigate the need for scripting in a recovery architecture. Specifically, with BIG-IP GTM, the scripting of DNS server changes can be eliminated and post-move IP address changes can be automated. With BIG-IP LTM, the need to reassign IP addresses can often be avoided altogether.

In this deployment guide we provide examples on how to address these two issues with SRM. In the first section, we show how BIG-IP LTM can be used to control IP addresses of guests between data centers. In the second section, we show how BIG-IP GTM can be used to deliver new DNS entries to client requests for the VM guests and/or Virtual IPs after the failover event. In the third section, we give an example on how to properly configure SRM applications using Access Policy Manager, as there are special considerations when using APM with route domains.

For this deployment, BIG-IP devices are required in the Protected Site and the Recovery Site. These BIG-IP devices service local traffic requests using LTM (with optional APM) and global traffic requests using GTM. In this scenario, the BIG-IP system does not play a role in storage replication; this is managed by specific storage replication adapters supported and distributed by storage vendors and VMware.

For more information on the F5 devices in this guide, see *http://www.f5.com/products/big-ip/*

Products and versions tested

Product	Version
BIG-IP LTM, GTM, APM	11.4 - 11.6
VMware vCenter SRM	5.8.0
VMware vSphere Replication	5.8.0.0
VMware ESXi server	5.5.0 Update 2 Enterprise Plus
VMware vCenter Server	5.5.0 Update 2 Enterprise Plus

Important: *Make sure you are using the most recent version of this deployment guide, available at http://www.f5.com/pdf/deployment-guides/vmware-site-recovery-manager-dg.pdf*

To leave feedback for this or other F5 solution documents, email us at *solutionsfeedback@f5.com*

Contents

Prerequisites and configuration notes

The following are general prerequisites and configuration notes for this guide:

- You must have BIG-IP Systems running LTM (and optionally APM) in both data centers to take advantage of Route Domain functionality.
- You must have at least one BIG-IP GTM pair to service DNS requests. Because SRM scenarios may require the shut down of an entire data center, we recommend GTM pairs in each data center.
- You must have a working installation of VMware vSphere with VMware SRM
- In this document, hosts refer to ESX Server Hosts. Guests or VMs refer to Virtual Machines.

More information

For more information on the topics discussed in this deployment guide, see the BIG-IP GTM Datasheet, available at *http://www.f5.com/pdf/products/big-ip-global-traffic-manager-ds.pdf*

Deployment overview

The configuration described in this guide uses the following workflow:

Configuration:

1. Configure Route Domains on the BIG-IP system
2. If using APM, configure an Access policy to include route domains.
3. Configure virtual machine guest networking to go through the BIG-IP system
4. Configure GTM to react to failover events either through health monitoring or iControl API automation

Failover or failback event:

1. Virtual machines are shut down by VMware
2. Storage is synchronized to the recovery site
3. Virtual machines are restored in the recovery site
4. GTM begins sending traffic to the new site

Preserving VM IP addressing during SRM recovery events

In this section, we begin by demonstrating how to use BIG-IP Route Domains to preserve IP addresses between data centers. VMware SRM is designed to shut down, move, and then bring up your hosts in the event of a planned site recovery event. Because the copy between sites is not live migration (the machines are shut down), stretching layer 2 traffic between the two data centers is not necessary, but a layer 2 data center interconnect may already be in place, complicating failover issues.

To solve the issue of IP addressing during SRM recovery events, VMware provides a scripted methodology which allows VMs to be reconfigured after they are restarted in the recovery site. This methodology has some drawbacks (discussed in detail below).

As an alternative, F5 proposes a solution using the Route Domain feature on the BIG-IP system. BIG-IP Route Domains enable administrators to isolate network traffic on the network. Route Domains allow sites to use the same IP address and subnet in more than one part of the network. Thus, the primary site and the secondary site can have the same IP address for the guests on the VMware hypervisor and use BIG-IP to shield the rest of the network from possible IP address conflicts.

This configuration is useful to sites that either:

- Have layer 2 stretched across the two data centers
- Cannot provision the same IP address because that range is taken or not available

In this scenario, all VMs participating in SRM ***must*** be located behind a BIG-IP device that supports Route Domains. It is also strongly suggested that the default route of the VMs is the BIG-IP for maximum protection against possible routing loops or other multiple IP address conflict scenarios. The VMware ESX servers themselves can be in any network and do not need to be behind the BIG-IP. Similarly, the storage networks need not be behind the BIG-IP system .

The connection of VM guests to the BIG-IP system can be achieved in a number of ways. Most simply, one NIC (or two NICs in a NIC teaming combination) on each ESX host can be physically connected directly to the BIG-IP. A slightly more scalable solution might be to connect the ESX servers to a switch and then to the BIG-IP (this approach might be preferable if there are more ESX hosts than available ports on BIG-IP). In either case, the goal of connectivity for VM guests is to shield them from all other traffic in the network and have VM guests be reachable only through the BIG-IP. If switches are used, care should be taken to avoid bridge loops or multiple ingress points to the VMs.

To implement Route Domains, a unique Route Domain identification number is used. That ID is appended to Virtual Servers, Pools, and other elements of a BIG-IP configuration. Using this configuration, it is even possible to have two hosts with the exact same IP address in the same data center.

The format for specifying a route domain object IP address is **A.B.C.D%ID**.

In this guide, we demonstrate how to setup route domains on each BIG-IP in your data center.

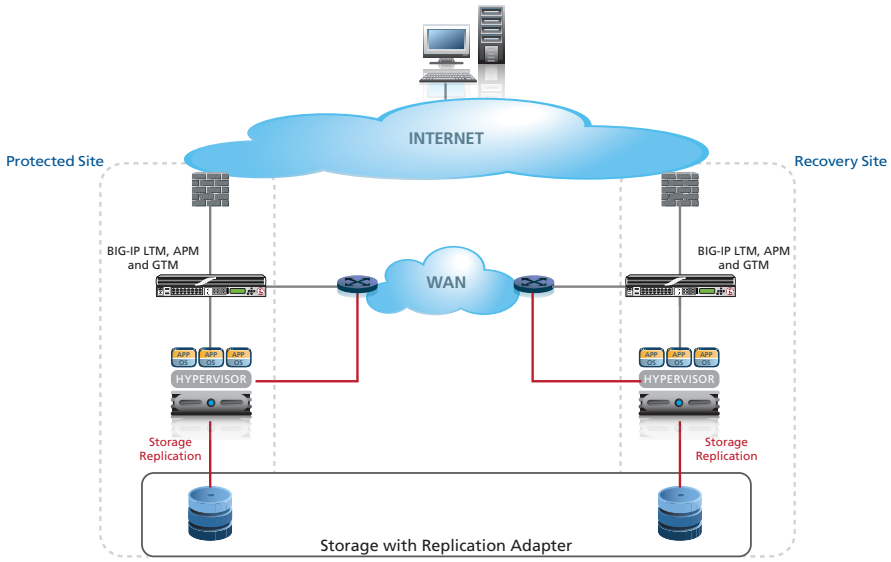

Figure 1: *Configuration example*

Note: *In this diagram, there may be a layer 2 switch between ESX hosts and the BIG-IP system as described in this document. Also, storage replication does not run through BIG-IP as indicated by the red line.*

Example configuration of BIG-IP Route Domains with VMware SRM

In this example we have two sites, a protected site in our primary data center, and a recovery site in our secondary data center. The IP address scheme on our primary data center is shown in the following tables.

Primary Data Center - Protected Site - without BIG-IP

Location of Host	Host Name	IP Address	Notes
Incoming from Internet	www.example.com	192.0.32.10	Fronts www.example.com web servers
Protected VMs	webserver1	10.0.1.200/24	
Protected VMs	webserver2	10.0.1.201/24	
Protected VMs	webserver3	10.0.1.202/24	
Incoming from Web Servers	apps.example.com	10.0.2.100	Load balances connections between web servers and application servers
Protected VMs	appserver1	10.0.2.200/24	
Protected VMs	appserver2	10.0.2.201/24	
Protected VMs	appserver3	10.0.2.202/24	

Secondary Data Center - Recovery Site - without BIG-IP

Location of Host	Host Name	IP Address	Notes
Incoming from Internet	www.example.com	65.61.115.222	Fronts www.example.com web servers
Protected VMs	webserver1	172.16.1.10/24	
Protected VMs	webserver2	172.16.1.11/24	
Protected VMs	webserver3	172.16.1.12/24	
Incoming from Web Servers	apps.example.com	172.18.1.100	Load balances connections between web servers and application servers
Protected VMs	appserver1	172.18.1.10/24	
Protected VMs	appserver2	172.18.1.11/24	
Protected VMs	appserver3	172.18.1.12/24	

In our example we are solving the issue of the 10.0**.1**.0/24 (serving web servers) network and the 10.0**.2**.0/24 (serving application servers) network not being available in the Recovery Site. There are two primary options to solving this issue with working with VMware SRM:

- Provision the 10.0.1.0/24 space and the 10.0.2.0/24 space in the Recovery Site.
- Use SRM in conjunction with VMware scripts to run the customization wizard (e.g., **sysprep** for windows) to change the IP address, gateway, and so on.

With the first option, provisioning the 10.0.1.0 and 10.0.2.0 network address may not be feasible either because of conflicts or because of routing loop concerns. For example, if there is an existing layer 2 data center interconnect, having the identical IP addresses available in both data centers would not be allowed.

The second option also presents potential problems. While sysprep works for Windows guests, it may not work for all Linux variants and other operating systems. Further, the customization of additional routes may not be possible. Finally, licensing and configuration issues may prevent an existing guest from being reassigned an IP address.

We are providing an alternative method using the BIG-IP; specifically the provisioning of the 10.0.1.0/24 and 10.0.2.0/24 space in a sandbox completely isolated by the BIG-IP system. This approach has several benefits and drawbacks that should be considered:

1. The first benefit is that BIG-IP provides strict isolation and even in networks with layer 2 data center interconnects will not cause conflicts or routing loops.
2. The second benefit is that this approach will work for all guests regardless of whether VMware supports syspep like functionality for the particular operating system.
3. The primary drawback is that the BIG-IP must be used as the default route for all connections into and out of the Virtual Machine. This means that BIG-IP should participate in routing.

With BIG-IP Route Domains configured, the following tables show what the IP addresses look like from the perspective of BIG-IP.

Primary Data Center - Protected Site with BIG-IP

Location of Host	Host Name	IP Address	Notes
Incoming from Internet	www.example.com	192.0.32.10	Fronts www.example.com web servers
Protected VMs	webserver1	10.0.1.200%10/24	
Protected VMs	webserver2	10.0.1.201%10/24	
Protected VMs	webserver3	10.0.1.202%10/24	
Incoming from Web Servers	apps.example.com	10.0.2.100%10	Load balances connections between web servers and application servers
Protected VMs	appserver1	10.0.2.200%10/24	
Protected VMs	appserver2	10.0.2.201%10/24	
Protected VMs	appserver3	10.0.2.202%10/24	

Secondary Data Center - Recovery Site with BIG-IP

Location of Host	Host Name	IP Address	Notes
Incoming from Internet	www.example.com	65.61.115.222	Fronts www.example.com web servers
Protected VMs	webserver1	10.0.1.200%10/24	
Protected VMs	webserver2	10.0.1.201%10/24	
Protected VMs	webserver3	10.0.1.202%10/24	
Incoming from Web Servers	apps.example.com	10.0.2.100%10	Load balances connections between web servers and application servers
Protected VMs	appserver1	10.0.2.200%10/24	
Protected VMs	appserver2	10.0.2.201%10/24	
Protected VMs	appserver3	10.0.2.202%10/24	

Implementing Route Domains

A route domain is a configuration object that isolates network traffic for a particular application on the network, allowing you to assign the same IP address or subnet to multiple nodes on a network, provided that each instance of the IP address resides in a separate routing domain. For more specific information on route domains, see the BIG-IP system documentation.
In order to implement Route Domains, you need to complete the following tasks:

1. Review and implement F5 guidance on routing traffic through your BIG-IP system. Remember that your guests (VMs) need to be able to reach all of their resources, such as DNS, NTP, etc through the BIG-IP system. See the following sections of the BIG-IP manuals for more information:
   https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/tmos-routing-administration-11-6-0.html?sr=44208215
   and
   https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/tmos-routing-administration-11-6-0/10.html?sr=44208239
2. Implement BIG-IP Route Domains (using the procedures in the following section).
3. Modify your switch fabric to provide connecting from your ESX hosts (and by reference your Virtual Machines) to BIG-IP. For specific information on modifying the switch fabric, refer to the switch documentation.
4. Modify your VM guests default route to point to the BIG-IP system. See the VMware documentation modifying VM hosts.

Configuring BIG-IP LTM

The first task is to configure route domains on the BIG-IP system. To configure Route Domains, you must create a VLAN, the Route Domain, and a self IP address. This configuration also includes a pool and virtual server. For more information on Route Domains, see *https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/tmos-routing-administration-11-6-0/8.html?sr=44208383*
Use the following table for guidance on configuring the BIG-IP LTM. The table contains a list of configuration objects along with any non-default settings you should configure as a part of this deployment. Settings not mentioned in the table can be configured as applicable for your configuration. For specific instructions on configuring individual objects, see the online help or product manuals.

VLANs (*Main tab > Network > VLANs*)	
Name	Type a unique name, such as **srm-vlan**
Tag	You can optionally assign a tag. If you do not, the system automatically assigns a tag.
Interfaces	Select the interface that has access to tagged traffic and then click Add.
Route Domain (*Main tab > Network > Route Domain*)	
Name	Type a unique name **All other settings are optional**
ID	Type a number for the Route Domain ID, using a number from 1 to 65534.
Strict Isolation	Check the box to **Enable** Strict Isolation.
VLANs Available	Select the VLAN you created and move it to the Members box.
Self IP address (*Main tab > Network > Self IPs*)	
Name	Type a unique name
IP Address	Type an IP address using the following syntax: **<IP address>%<Route Domain ID>**. For example: 192.0.32.254%10.
Netmask	Type the corresponding subnet mask. For example: 255.255.255.0.
VLAN	Select the VLAN you created and move it to the Members box.
Pools (*Main tab > Local Traffic > Pools*)	
Name	Type a unique name
Load Balancing Method	Choose your preferred load balancing method. We use **Least Connections (Member)**.
Address	Type the appropriate IP address of one of the Virtual Machines, followed by %10. For example, **10.0.1.200%10**.
Service Port	Type the appropriate service port. If this VM is not participating in application delivery controller functionality such as load balancing, you can use the wildcard service port (**0**) to make all ports available. Click **Add** to repeat Address and Port for all VMs.
SNAT Pool[1] (*Main tab > Local Traffic > Address Translation > SNAT Pool List*)	
Name	Type a unique name
IP Address	Type an otherwise unused IP address for the SNAT pool[1], and then click **Add**. Repeat for any additional addresses needed.

[1] SNAT pools are helpful for pool members with no virtual servers attached to them in this deployment example. SNAT pools will allow these pool members to communicate while using the BIG-IP as their translation source. For pool members with virtual servers, or in your specific deployment, SNAT pools may be optional or not needed. For more information on SNAT pools, see the Configuring SNATs chapter in the Configuration Guide for Local Traffic Management.

Virtual Servers (*Main tab > Local Traffic > Virtual Servers*)	
Name	Type a unique name
Destination Address	Type the IP address you want to use for this virtual server
Service Port	Type the appropriate service port
Default Pool	Select the pool you created
If applicable, you can select the appropriate Profiles required for your application.	

Repeating the configuration on the BIG-IP LTM in the secondary data center

After completing the configuration on the BIG-IP LTM in the primary data center, the next task is to repeat this entire section of the deployment guide on the BIG-IP LTM in the secondary data center.

For ease of management in our example, we use the same route domain in both data centers. This is optional, and any route domain ID may be used of your choosing.

Return to the start of the table and repeat the configuration objects on the secondary BIG-IP LTM.

➲ **Important:** *You must create all of the objects on the BIG-IP in the secondary data center for this deployment to function properly*

Special considerations for configuring the BIG-IP system for using APM and Route Domains

If you want to use route domains in your implementation along with BIG-IP Access Policy Manager (APM), you must use the following guidance to configure the BIG-IP system.

To configure the BIG-IP system for APM and route domains

1. Create a new partition on the BIG-IP system (click **System > Users > Partition List > Create**).
2. Create a new route domain and make it default for your new partition (click **Network > Route Domains > Create**).
3. Switch to your new partition (the Partition list is in the upper right corner of the Configuration utility) and create a new VLAN, Self IP, and Route (if applicable) in the new partition.
4. While still in the partition you created, create an Access Policy along with appropriate policy objects (**Access Policy > Access Profiles > Create**). Alternatively, you could use an iApp template to configure the BIG-IP APM Access Policy and associated objects, if applicable for your application.

See *https://support.f5.com/kb/en-us/products/big-ip_apm/manuals/product/apm-implementations-11-6-0/3.html#unique_1025147473* for additional information on routing with Access Policies.

Switching site availability in SRM using BIG-IP GTM

With BIG-IP Global Traffic Manager (GTM), DNS requests for an entire data center can be redirected to a new location based either on monitors or scripting. In this SRM deployment, GTM's control of SRM guests can be straightforward because entire data centers are usually vacated during failover events. However, GTM can also easily handle more complex deployments and can direct traffic to two data centers in an active-active scenario, using various algorithms to determine the ratio of requests.

With GTM, DNS can be deployed in one of two ways. First, either GTM can front-end requests for a BIND server, or GTM can take over all DNS requests. For specifics on deploying GTM, please see the BIG-IP GTM documentation. Next, GTM actively monitors BIG-IP LTM looking for triggers that indicate the data center is no longer available. In the case of recovery events, GTM will start to direct traffic to the second data center after the last virtual machine has shut down.

Another consideration in this deployment is whether GTM will be serving traffic both internally and externally. If so, care should be taken to create a DNS architecture that serves all of the needs of your enterprise. The configuration of GTM for DNS is outside the scope of this document.

Depending on the desired level of automation of control, there are two ways to deploy GTM to react to failover events. For complete automation, the GTM can be configured to automatically discover a failover event using monitors and react to it accordingly. If less automation is desired, it is possible to use SRM to initiate the failover event on the GTM for the DNS addresses given to clients for an associated pool. This method is configured using GTM's iControl API.

GTM configuration overview

The BIG-IP GTM can be configured to direct traffic in each data center in a number of ways, depending on if your architecture is active/active or active/standby.

- **Active/Active**
  If your data center architecture can accommodate traffic incoming and outgoing to both data centers at the same time (active/active), then GTM can be setup with ratio based global load balancing. In this case, the ratio of DNS entries delivered to incoming client DNS requests is automatically adjusted based on the number of active VMs in each date center.

  For example, if three VMs are active in the primary data center and seven VMs are active in the recovery data center, GTM provides the DNS entry for the recovery data center 70 percent of the time.

- **Active/Standby**
  In the active/standby configuration, GTM only delivers the DNS entry of the active data center to incoming client DNS requests.

Setting up the BIG-IP GTM for failover is a multi-step process:

- Assign all Virtual Machines to pools on the BIG-IP LTM
- Use GTM to monitor these pools and provide the IP addresses to external clients based on the state of the pool members. When GTM detects the pool is down, it starts providing users with the recovery site IP addresses.

➲ **Note**: *Virtual Servers are not required for all pool members. SNATs may be used where a pool member needs to communicate to the outside network but does not necessarily need a Virtual Server for application delivery purposes.*

Configuring the BIG-IP GTM for automatic switching of site availability

Next, we configure the BIG-IP Global Traffic Manager for switching site availability in SRM.

Configuring the BIG-IP GTM

Use the following procedures to configure the BIG-IP Global Traffic Manager for Global Server Load Balancing using the VS Score load balancing method. For specific instructions on configuring individual objects, see the online help available from the Help tab, or the BIG-IP GTM documentation.

GTM Object	Description/Notes	
Listener *(Global Traffic -->Listeners)*	***Name***	Type a unique name
	Destination	Type the IP address on which the Global Traffic Manager listens for network traffic. In our example, this is an IP address on the WAN network.
	VLAN Traffic	Select a VLAN setting appropriate for this Listener.
	Create additional listeners using the same IP address if necessary. If creating an IPv6 listener, be sure to use an IPv6 destination address	
Data Center *(Global Traffic -->Data Centers)*	***Name***	Type a unique name. Configure other options as applicable for your environment. Repeat for the secondary data center.
Servers *(Main tab-->Global Traffic -->Servers)*	***Name***	Type a unique name
	Product	Select the either **BIG-IP System (Single)** or **BIG-IP System (Redundant)**. Redundant is only used when the GTM is also an LTM/GTM combo and specifically configured for LTM failover of the listener. Otherwise use **BIG-IP System (Single**).
	Address List: Address	Type the Self IP address of this GTM.
	Data Center	Select the Data Center you created
	Health monitors	*Optional:* Select **bigip**
	Virtual Server Discovery	**Enabled** (We strongly recommend Enabling Discovery, however you can leave this set to Disabled and manually configure the virtual server information)
	Repeat 2 more times for the Secondary LTM and the Primary GTM. When you are finished you should have three GTM Server objects, one referencing the GTM itself, one referencing the Primary site's LTM, and one referencing the secondarily site's LTM.	
Enabling connectivity with remote BIG-IP systems *(Command line)*	When adding a remote BIG-IP LTM server, you must make sure the ***big3d*** agent is on the same version on the BIG-IP LTM. If you have never registered the BIG-IP LTM systems with BIG-IP GTM before, you should perform the following steps from GTM using the management IP address(es) of each of the LTM hosts. From the GTM device command line, type: **big3d_install** ***<IP address of target system>*** where the target system is the BIG-IP APM that you want to add as a server on the GTM. This pushes out the newest version of big3d. Next, type: **bigip_add** to exchange SSL keys with the BIG-IP LTM. Type the password at the prompt, and then type **iqdump** ***<ip address of remote box>***. If the boxes are communicating over iQuery, you see a list of configuration information from the remote BIG-IP. The **bigip_add** command must be run for every BIG-IP in the configuration. Return to your gtm administrative console and you should now note your Primary and Secondary BIG-IP servers with a healthy status displayed as a green circle.	
Pools *(Global Traffic -->Wide IPs --> Pools)*	***Name***	Type a unique name
	Load Balancing Method	***Preferred:*** **VS Score**[1] (if using Topology-based GTM configuration, select **Topology** here) ***Alternate:*** **VS Capacity** ***Return to DNS:*** **VS Score**
	Member List	**Virtual Server** Select the appropriate virtual server you created for the application from the dynamically populated list and then click **Add**. You must select the virtual server by IP address and port number. Repeat process and select secondary SRM site virtual server by IP and port number.
Wide IPs *(Global Traffic -->Wide IPs)*	***Name***	Type the FQDN for your application.
	Load Balancing Method	**Topology**
	Pool List	Select the pool you created, and then click Add. Repeat for the secondary Pool.

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

Conclusion

In this guide, we have shown options for deploying VMware SRM with the BIG-IP LTM, GTM, and APM. Preserving IP addresses with the help of LTM route domains allows administration to sandbox IP address space for virtual machines participating in SRM even if layer 2 data center interconnections are in place. With GTM, either configured for an active/active or active/passive data center, full administrative control over DNS is provided. Administrators should see a reduction in scripting required around SRM. This solution is compatible with APM and outlines the preferred configuration when using APM with route domains.

Document Revision History

Version	Description	Date
1.0	New guide for BIG-IP v11.4	03-27-2015

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com