



Deploying the BIG-IP System with VMware vCenter Site Recovery Manager

Contents

- 2 Prerequisites and configuration notes
- 2 Deployment overview
- 3 Example configuration of BIG-IP Route Domains with VMware SRM
- 6 Implementing Route Domains
- 9 Switching site availability in SRM using BIG-IP GTM

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) and Global Traffic Manager (GTM) 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.

For this deployment, BIG-IP devices are required in the Protected Site and the Recovery Site. These BIG-IPs service local traffic requests using LTM 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	10.2.1
VMware SRM	4.1
VMware vSphere	vSphere 4 Enterprise Plus
VMware ESX server	4.1
VMware vCenter Server	4.1

Document Version

1.0

Prerequisites and configuration notes

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

- You must have BIG-IP Systems running LTM 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.

Deployment overview

The configuration described in this guide uses the following workflow:

Configuration:

1. Configure Route Domains on BIG-IP
2. Configure virtual machine guest networking to go through BIG-IP
3. Configure GTM to react to failover events either through health monitoring or iControl API automation

Failover or fallback 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.

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.

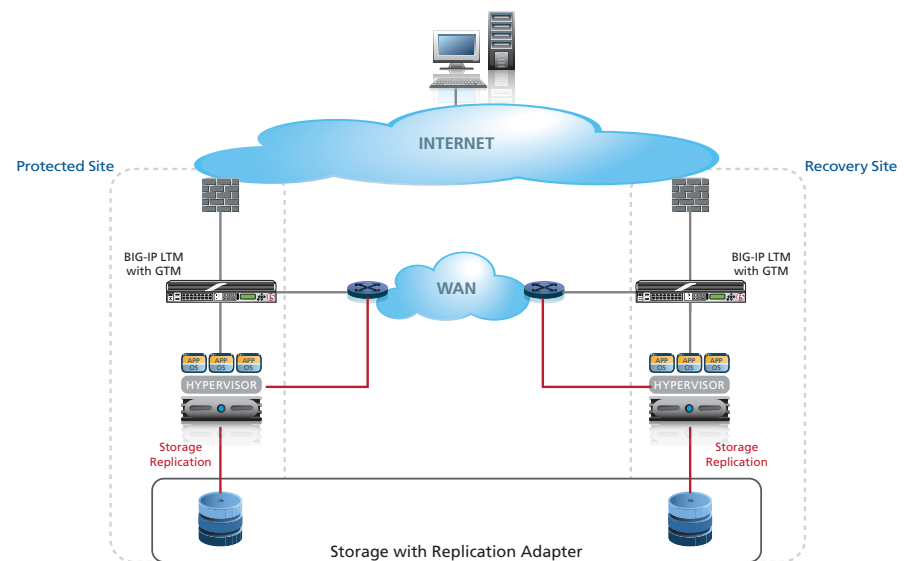


Figure 1: Configuration example

➔ **Note** In the above diagram, there may be a layer 2 switch between ESX hosts and the BIG-IP 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 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 sysprep 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

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. Remember that your guests (VMs) need to be able to reach all of their resources, such as DNS, NTP, etc through the BIG-IP. See the following sections of the BIG-IP manuals for more information:
http://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/tmos_management_guide_10_1/tmos_routes.html#1006923
 and
http://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/tmos_management_guide_10_1/tmos_routes_dynamic.html#1192201
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 BIG-IP. For specific information on modifying your VM hosts, see the VMware documentation.

Configuring BIG-IP Route Domains

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.

Creating a VLAN

The first task is to create a VLAN

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 we use **srm-vlan**.
4. In the **Tag** box, you can optionally type a tag. In our example, we leave this blank. The BIG-IP 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 appropriate box by clicking the Add (<<) button. In our example, we select 1.14.
6. Click the **Finished** button.

Creating the Route Domain

The next task is to create the Route Domain.

To create a Route Domain

1. On the Main tab, expand **Network** and then click **Route Domains**.
2. Click the **Create** button.
3. In the **ID** box, type a number for the Route Domain ID, using a number from 1 to 65534. We use an identifying network tag, and type **%10**.
4. *Optional:* In the **Description** box, type a description.
5. Click the **Strict Isolation** box to enable Strict Isolation.
Strict isolation controls whether cross-routing restrictions are enabled. For example, if there are other route domains on the BIG-IP, this setting will control if one route domain can communicate with the other route domain on the same BIG-IP. In our example, we do not want any cross-routing and enable Strict Isolation.
6. From the **VLANs Available** list, select the VLAN you created in the preceding section and click the Add (<<) button to move it to the selected box.
7. Leave all other settings at the default levels.
8. Click **Finished**.

Creating a self IP address

The next task is to create a self IP address on the BIG-IP LTM that uses Route Domains.

To create a self IP address

1. On the Main tab, expand **Network**, and then click **Self IPs**.
2. Click the **Create** button. The new Self IP screen opens.
3. In the **IP Address** box, type an IP address using the following syntax:
<IP address>%<Route Domain ID>
In our example, we type **192.0.32.254%10**.

4. In the **Netmask** box, type the corresponding subnet mask.
In our example, we use 255.255.255.0.
5. From the **VLAN** list, select the VLAN you created in *Creating a VLAN on page 7*.
In our example, we select **srm-vlan**.
6. Click the **Finished** button.

Creating a pool

The next task is to create a load balancing pool on the BIG-IP system.

To create a pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**.
2. Click the **Create** button.
3. In the **Name** box, type a unique name.
4. In the Resources section, from the **Load Balancing Method** list, choose your preferred load balancing method.
5. In the New Members section, you add the virtual machines to the pool:
 - a. In the **Address** box, type the appropriate IP address of one of the Virtual Machines, followed by **%10**. For example, **10.0.1.200%10**.
 - b. In the **Service Port** box, type the appropriate service port number for this device.
If this VM is not participating in application delivery controller functionality such load balancing, then you can use the wildcard service port to make all ports available.
 - c. Click the **Add** button to add the member to the list.
 - d. Repeat steps a-c for each VM you want to add to the pool.
6. Click the **Finished** button.

Creating a SNAT pool

In this procedure, we create a SNAT pool. 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.

To create the SNAT pool

1. On the Main tab, expand **Local Traffic**, and then click **SNATs**.
2. On the Menu Bar, click **SNAT Pool List**.
3. Click the **Create** button.
4. In the **Name** box, type a name for this SNAT pool. In our example, we type **srm-snat-pool**.
5. In the **IP Address** box, type an otherwise unused IP address, and click the **Add** button.
6. Repeat this step for any additional address needed.
7. Click **Finished**.

Creating a virtual server

The final task in this section is to create a virtual server on the BIG-IP LTM.

To create the virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this virtual server.
4. In the Destination section, select the **Host** option button.
5. In the **Address** box, type the IP address of this virtual server. In our example, we type **192.0.32.100**.
6. In the **Service Port** box, type the appropriate service port.
7. From the **SNAT Pool** list, select the name of the SNAT Pool you created in the preceding procedure. In our example, we select **srm-snat-pool**.
8. In the Resources section, from the **Default Pool** list, select the pool you created.
9. Click the **Finished** button.

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 *Creating a VLAN on page 7* and repeat each procedure on the secondary BIG-IP LTM.

➔ **Important:** *You must perform all of the procedures on the BIG-IP in the secondary data center for this deployment to function properly*

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 data 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.

Creating the data centers

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

To create the data centers

1. On the Main tab, expand **Global Traffic** and then click **Data Centers**.
2. Click the **Create** button. The New Data Center screen opens.

3. In the **Name** box, type a name. In our example, we type **Primary**.
4. Complete the rest of the configuration as applicable for your deployment.
5. Click the **Finished** button.
6. Repeat this procedure for the Secondary data center.

Creating the monitor

The next step is to create a health monitor. The monitor type (in step 4) depends on the type of application you are running. In the following example, we create an HTTP monitor.

To create the monitor

1. On the Main tab, 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. In our example, we type **SRM-http-monitor**.
4. From the **Type** list, select the appropriate monitor type. In our example, we type **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.

To create the GTM servers

1. On the Main tab of the navigation pane, expand **Global Traffic** and click **Servers**.
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 **Primary-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 11*, and click the Add (<<) button. In our example, we select **SRM-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 primary data center, and one that includes the LTM virtual server in the secondary data center. The secondary 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**.
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 **Primary_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 11*, and click the Add (<<) button. In our example, we select **SRM-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. You must select the virtual server by IP address and port number. 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.

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. In our example, we type **www.example.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 11*, and then click the **Add** button. In our example, we select **Primary_pool**.
7. Repeat this step for the remote pool.
8. All other settings are optional, configure as appropriate for your deployment.
9. Click the **Finished** button.

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

Configuring the BIG-IP GTM for SRM-controlled switching of site availability

iControl is the BIG-IP API that allows scripting of BIG-IP modules including GTM. There are code samples on devcentral.f5.com that demonstrate how to use any of the various API languages to control the scripting of GTM. For more information on iControl and to download the SDK, see <http://devcentral.f5.com/icontrol>.

The following example leverages iControl and Perl to control GTM pool members. You can also use PowerShell (see <http://devcentral.f5.com/wiki/default.aspx/iControl/PowerShell/GTMPoolMember.html>) for the PowerShell code.

Perl

In the following sample for the perl based iControl API for GTM, several methods give SRM the ability to change the state of the pool based on the decision on the administrator versus the monitoring of GTM.

Enabling network connections for iControl involves allowing TCP port 443 access to the BIG-IP, and creating a user to allow login on the BIG-IP. On the SRM host, the iControl API should be installed, using the following iControl Quickstart Guide will be useful:
<http://devcentral.f5.com/Default.aspx?tabid=63&articleType=ArticleView&articleId=59>

Once communication is established between your SRM hosts and GTM, the Perl GTM Pool Member script may be used as a customizable template for your SRM environment. The full script can be found at
<http://devcentral.f5.com/wiki/default.aspx/iControl/PerlGtmPoolMember.html>

In this example, the usage is:

```
Usage: GTMPoolMember.pl host uid pwd [wideip [pool [member [enable|disable]]]]
```

SRM calls the perl script, in our example, with the following settings:

```
GTMPoolMember.pl bigip-gtm1 myuser mypassword 192.0.32.10 webserver 10.10.1.200 disable
GTMPoolMember.pl bigip-gtm1 myuser mypassword 192.0.32.10 webserver 10.10.1.201 disable
GTMPoolMember.pl bigip-gtm1 myuser mypassword 192.0.32.10 webserver 10.10.1.201 disable

GTMPoolMember.pl bigip-gtm1 myuser mypassword 65.61.115.222 webserver 10.10.1.200 enable
GTMPoolMember.pl bigip-gtm1 myuser mypassword 65.61.115.222 webserver 10.10.1.201 enable
GTMPoolMember.pl bigip-gtm1 myuser mypassword 65.61.115.222 webserver 10.10.1.201 enable
```

This disables the web server in the protected data center and enables the web servers in the recovery data center. Multiple calls to the script are necessary to complete the failover.

Conclusion

In this guide, we have shown options for deploying VMware SRM with the BIG-IP LTM and GTM. 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 datacenter, full administrative control over DNS is provided. Administrators should see a reduction in scripting required around SRM.

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>

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

Document Revision History

Version	Description
1.0	New Version

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

