Deployment Guide

Deploying F5 BIG-IP Global Traffic Manager™ on VMware™ vCloud® Hybrid Service™

A. Introduction

VMware vCloud Hybrid Service is an effective, flexible and reliable platform for enterprise customers to move application workloads to the cloud. The recent introduction of the Disaster Recovery as a Service offering from VMware extends the flexibility of the vCloud Hybrid Service platform by enabling customers to adopt Disaster Recovery and Business Continuity best practices.

F5 BIG-IP Global Traffic Manager™ (GTM) can play a crucial role in any disaster recovery strategy by allowing for seamless application failover between sites, ensuring that application availability is maintained throughout the DR process. BIG-IP GTM allows DNS Global Server Load Balancing (GSLB) to redirect user connections to applications running in vCloud Hybrid Service, without interruption.

1. BIG-IP GTM and BIG-IP Local Traffic Manager Benefits

F5 Big-IP GTM can be configured as a full proxy for GSLB applications and DNS across architectures – and across the globe. For greater flexibility, you can use BIG-IP Global Traffic Manager Virtual Edition (VE) to extend DNS services and global application availability to cloud or virtual environments, and maintain centralized control within the data center. Big-IP GTM on vCloud Hybrid Service offers:

- Seamless application failover and business continuity between the enterprise data center and vCloud Hybrid Service
- Global availability of applications between data centers and multiple vCloud Hybrid Service locations
- High performance of applications, whether on-premises or in the VMware cloud.

Additionally, customers who deploy BIG-IP Local Traffic Manager to help deliver applications to users in a reliable, secure, and optimized way can evaluate and test DNS and global app delivery functionality by provisioning BIG-IP Global Traffic Manager.

2. Use Cases Validated During Testing Phase

This document outlines the steps to deploy F5 in a vCloud Hybrid Service environment. Typical use cases for BIG-IP Global Traffic Manager are as follows:

- Multi-site or region availability of applications with intelligent direction to most available or closest application resource
- Disaster Recovery and Business Continuity- provide seamless DNS failover and user redirection to application resources in the cloud
- DNS security and scale.

In this validation, we tested two separate use cases; the first use case was centered on failover of a web application to vCloud Hybrid Service – Disaster Recovery. In this scenario, BIG-IP Global Traffic Manager provides the DNS failover between the primary on-premises location and the vCloud Hybrid Service disaster recovery location.

In the second scenario we demonstrated BIG-IP GTM to provide global availability of a

Horizon\View deployment across two separate instances of vCloud Hybrid Service.

In addition to the two scenarios validated, BIG-IP GTM can provide DNS availability for numerous applications deployed in vCloud Hybrid Service.

In this document we tested and validated BIG-IP GTM to enable application failover utilizing VMware vCloud Hybrid Service - Disaster Recovery. We tested the following two scenarios:

1. Disaster Recovery\Business Continuity of a web-based application utilizing VMware vCloud Hybrid Service - Disaster Recovery and BIG-IP GTM.
2. Global availability and global load balancing of a VMware Horizon\View web based connection server utilizing vCloud Hybrid Service in multiple locations.

In both cases, BIG-IP Global Traffic Manager was a key enabling technology to facilitate failover between data center locations.

B. Use Case Architecture and Deployment

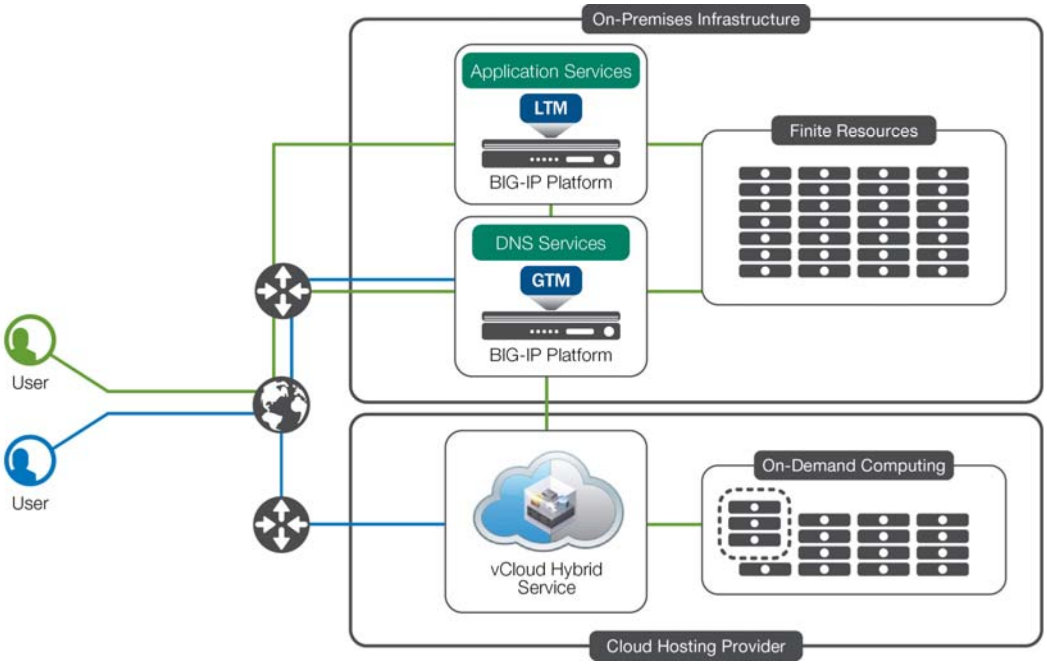

Figure 1: BIG-IP for Disaster Recovery and Multi-Site Environments

The first example involves set up of a WordPress configuration to simulate a packaged application server use case on vCloud Hybrid Service. For this example, we deployed two BIG-IP devices provisioned and licensed for GTM. The first GTM was deployed in the primary on-site location in Washington with a secondary GTM deployed as part of a synchronization group within vCloud Hybrid Service in Virginia. BIG-IP GTM Synchronization groups synchronize object configurations between GTM devices so that DNS responses to Wide-IPs can be maintained across locations.

Here are the steps to provision and deploy BIG-IP in vCloud Hybrid Service:

- Download the BIG-IP VE package *ovf from https://downloads.f5.com you will need to download the VMware VE 11.5.1 product
- Once the BIG-IP VE package is downloaded upload the BIG-IP VE into the vCloud Hybrid Service My Catalog
- In vCloud Hybrid Service click on 'Add Virtual Machine' select your resources and choose the 'My Catalog' tab
- Provide a name for your BIG-IP and ensure a public IP address is assigned to your primary management interface. You will need to setup firewall rules in vCloud Hybrid Service to provide access to the management IP address
- Once the BIG-IP is deployed navigate to https://bigippublicipaddress
- Use the default username: Admin password: Admin logon
- License your BIG-IP via the automatic method
- For additional details on deployment of BIG-IPVE please go to support.f5.com.

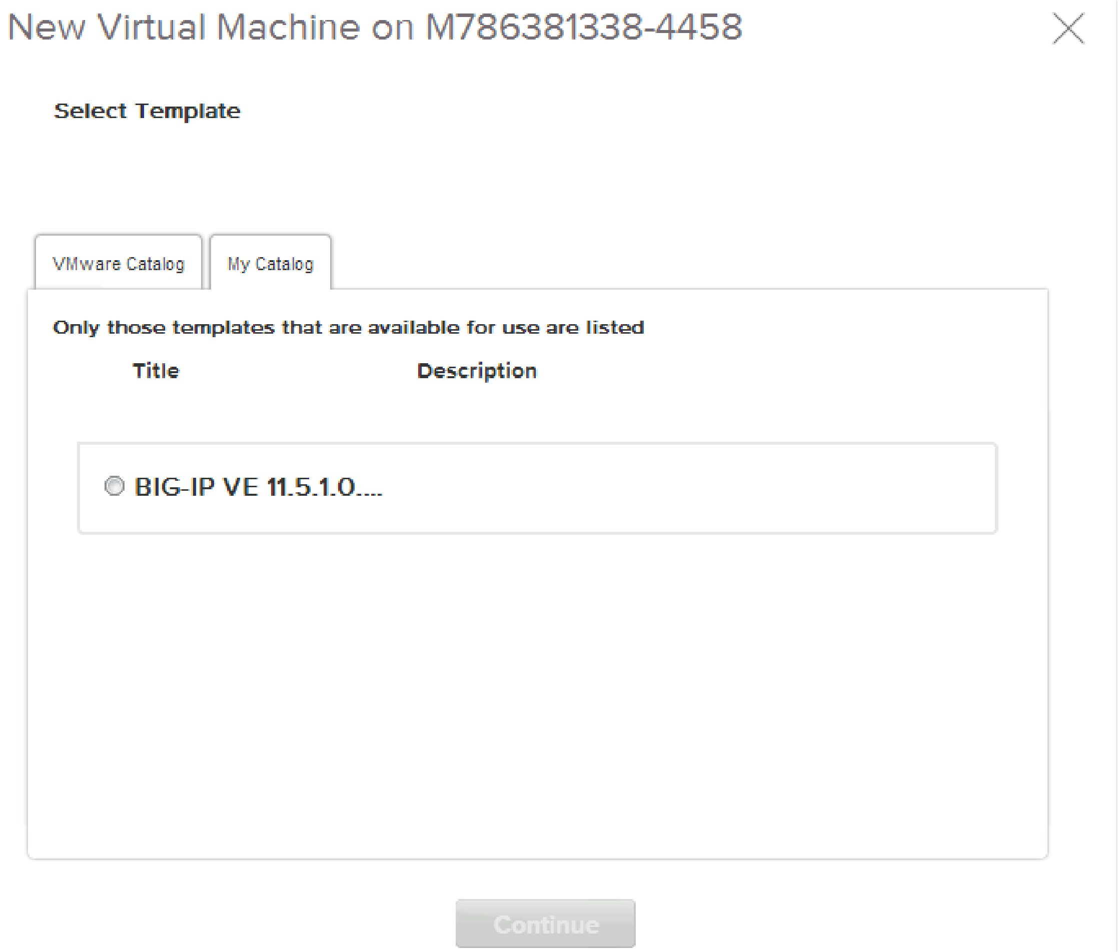
New Virtual Machine on M786381338-4458

Select Template

Only those templates that are available for use are listed

Title	Description
○ BIG-IP VE 11.5.1.0....	

1. Configure VLANs and Self IP Addresses:

In this deployment two VLANs were created for a simple two-armed deployment. An Internal VLAN and an External VLAN were created as part of the initial setup of the BIG-IP in each location as shown in Figure 2 below. After configuring the external and internal VLANs on each BIG-IP we assigned two Self-IP addresses for each VLAN, shown in Figure 3.

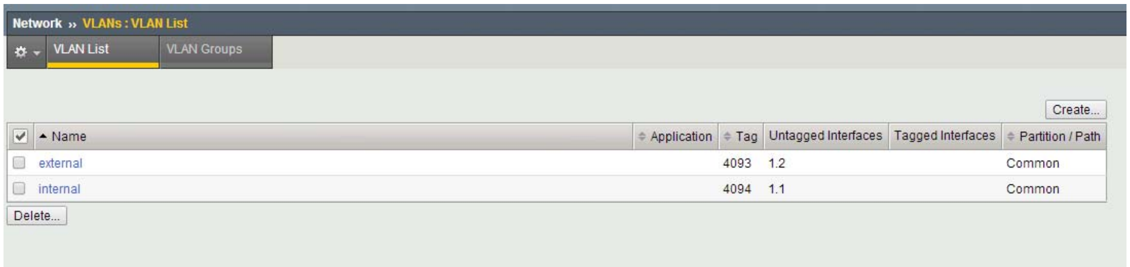

☑	Name	Application	Tag	Untagged Interfaces	Tagged Interfaces	Partition / Path
□	external		4093	1.2		Common
□	internal		4094	1.1		Common

Figure 2: External and Internal VLAN Setup Untagged Interface 1.1 (Internal) and 1.2 (External)

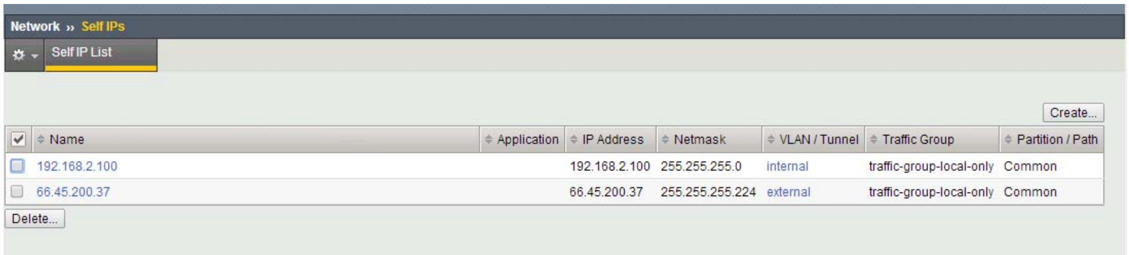

☑	Name	Application	IP Address	Netmask	VLAN / Tunnel	Traffic Group	Partition / Path
□	192.168.2.100		192.168.2.100	255.255.255.0	internal	traffic-group-local-only	Common
□	66.45.200.37		66.45.200.37	255.255.255.224	external	traffic-group-local-only	Common

Figure 3: Self IP Addresses on vmtm-BIGIP01

2. Create Data Centers

In order to setup and deploy BIG-IP, a data center object will need to be created for each site to which you want to load balance DNS requests. In the example below, we have created two data center objects ***DC_Primary*** and ***Failover****.* To create the data center objects expand **the *DNS* menu and then *GSLB ->Data Centers.*** *Click '***Create'** to configure the Data Center Objects.

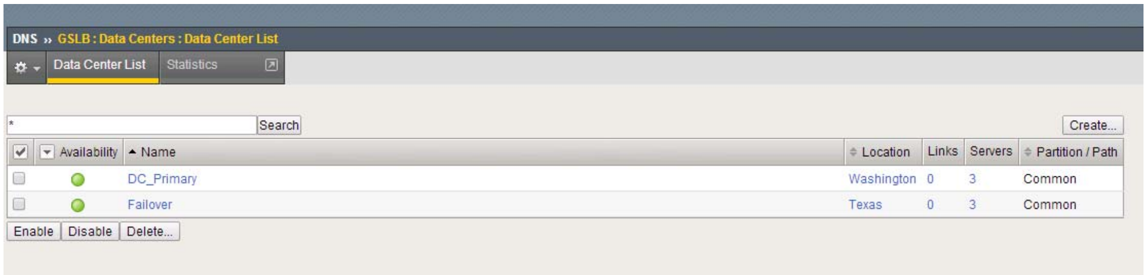

☑	Availability	Name	Location	Links	Servers	Partition / Path
□	●	DC_Primary	Washington	0	3	Common
□	●	Failover	Texas	0	3	Common

Figure 4: Data Center Object Creation

3. Create Servers

After creating the data center objects a server object will need to be created for each GTM and each server object in the configuration. In the below example we have created server objects for each GTM ***GTMA & GTMB*** *and a server object for the primary and secondary* ***WordPress*** *server.*

To create GTM server object expand the **DNS** menu and then navigate to **GSLB -> Servers**. From the screen shown in Figure 5, click ‘**Create.’**

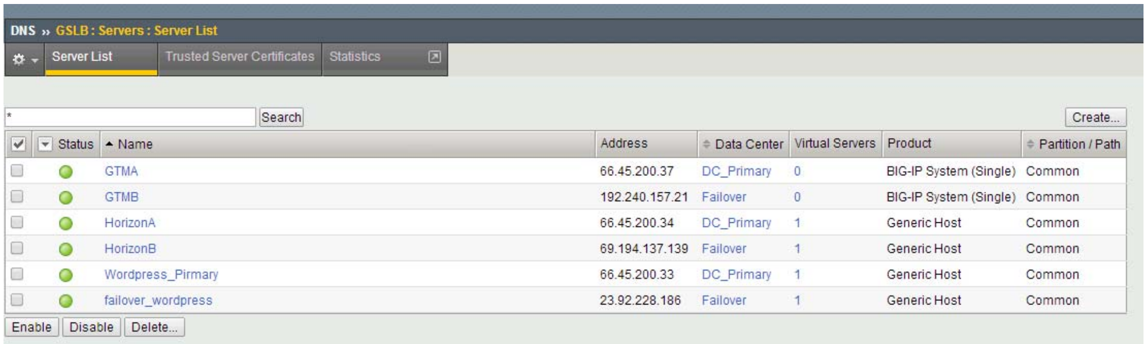

DNS » GSLB : Servers : Server List

☑	Status	Name	Address	Data Center	Virtual Servers	Product	Partition / Path
□	●	GTMA	66.45.200.37	DC_Primary	0	BIG-IP System (Single)	Common
□	●	GTMB	192.240.157.21	Failover	0	BIG-IP System (Single)	Common
□	●	HorizonA	66.45.200.34	DC_Primary	1	Generic Host	Common
□	●	HorizonB	69.194.137.139	Failover	1	Generic Host	Common
□	●	Wordpress_Pirmary	66.45.200.33	DC_Primary	1	Generic Host	Common
□	●	failover_wordpress	23.92.228.186	Failover	1	Generic Host	Common

Figure 5: Create Server Objects

After clicking on **Create** you will be prompted to complete a menu of questions. We used the following in our GTM configuration:

Name	GTMA
Product	Choose BIG-IP System Single
Address	Enter your Public IP and Translated Private IP if Applicable
Data Center	Choose Primary Data Center from Drop Down
Prober Pool	Accept Default
Status	Keep Default ‘Enabled’
Configuration	Choose ‘Advanced’ from drop down
Health Monitors	We selected gateway ICMP for basic availability

Accept all other option defaults and click ‘**Update’.** After creating the first GTM object create a second **GTMB** object. Select all of the above configuration options shown in Figure 6 but ensure that the IP address is a configured Self-IP on the secondary GTM.

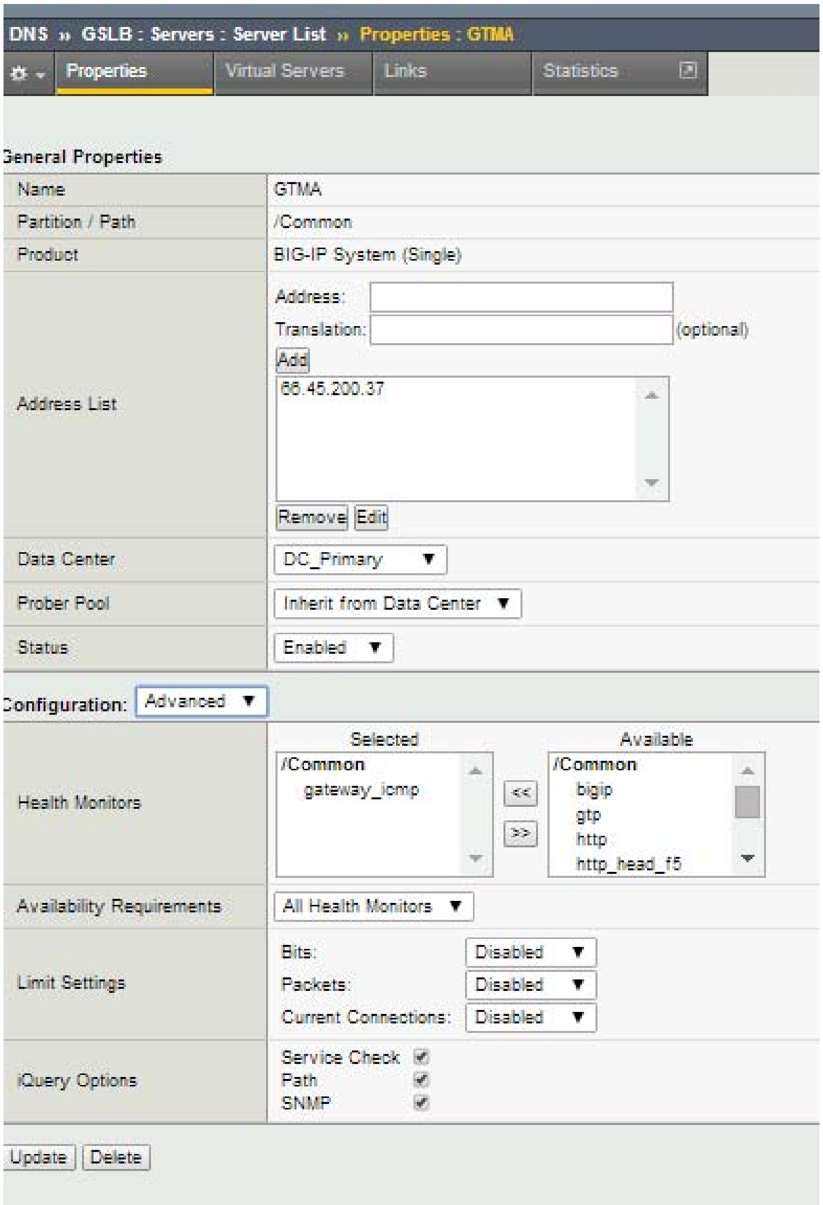

DNS » GSLB : Servers : Server List » Properties : GTMA

General Properties

Name	GTMA
Partition / Path	/Common
Product	BIG-IP System (Single)
Address List	Address: Translation: (optional) 66.45.200.37
Data Center	DC_Primary
Prober Pool	Inherit from Data Center
Status	Enabled

Configuration: Advanced

Health Monitors	Selected: /Common gateway_icmp	Available: /Common bigip gtp http http_head_f5
Availability Requirements	All Health Monitors	
Limit Settings	Bits: Disabled Packets: Disabled Current Connections: Disabled	
iQuery Options	Service Check ☑ Path ☑ SNMP ☑	

Figure 6: Create GTM Server Object

4. Create WordPress Server Object and Virtual Server

(Wordpress is used to simulate a packaged application. In a real world scenario, this use case will apply to any packaged application)

From the server configuration click create to configure the WordPress Server Objects:

Name	WordPress_Primary
Product	Choose ‘Generic Host’
Address & Translation	Enter Public IP and Translated Private IP Address

Data Center	DC_Primary
Prober Pool	Keep Defaults
Status	Keep Default ‘Enabled’
Configuration	Choose Drop Down and Select ‘Advanced’
Health Monitors	Choose ‘HTTPS’

Once you have configured the server object you will need to configure the Virtual Server object associated with the primary WordPress server. Enter the below items in the ‘**Resources’** section.

Virtual Server Discovery	Choose ‘Disabled’ from the drop down
Name	WordPress A
Address	Public IP Address
Translation	Private IP Address
Monitor	Choose HTTP

Once completed click ‘Update’.

Create another failover object for the failover **WordPress** server and **WordPress** Virtual Server using the above parameters. See diagram in Figure 7.

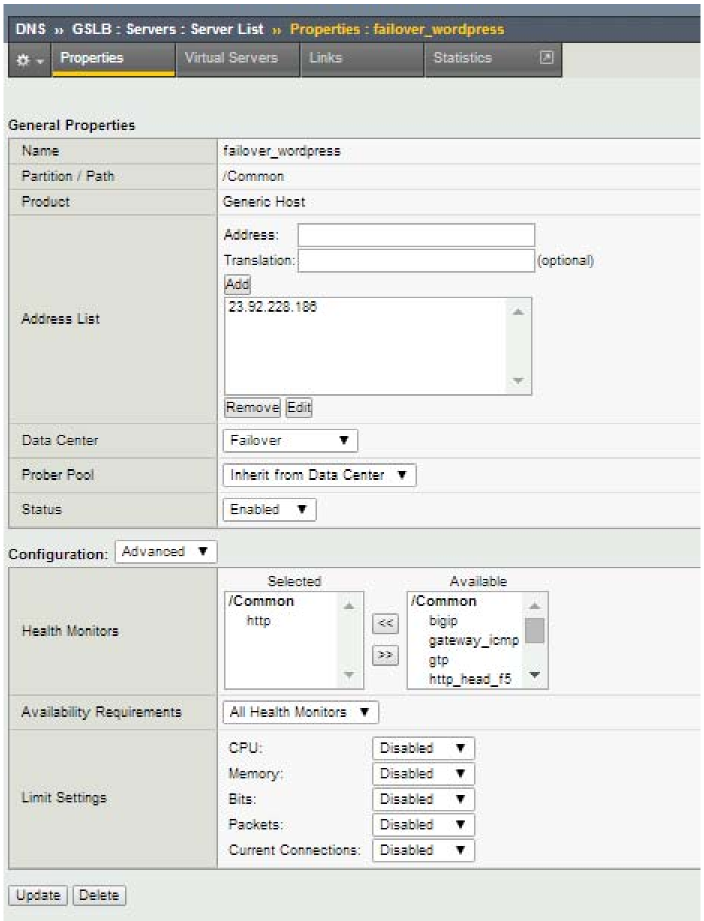

General Properties

Name	failover_wordpress
Partition / Path	/Common
Product	Generic Host
Address List	Address: Translation: (optional) 23.92.228.186
Data Center	Failover
Prober Pool	Inherit from Data Center
Status	Enabled

Configuration: Advanced

Health Monitors	Selected: /Common http Available: /Common bigip gateway_icmp gtp http_head_f5
Availability Requirements	All Health Monitors
Limit Settings	CPU: Disabled Memory: Disabled Bits: Disabled Packets: Disabled Current Connections: Disabled

Figure 7: Failover WordPress Server

Change the IP address and translation to the public IP and private IP address of the failover server. Assign this to the failover data center location as shown in Figure 8.

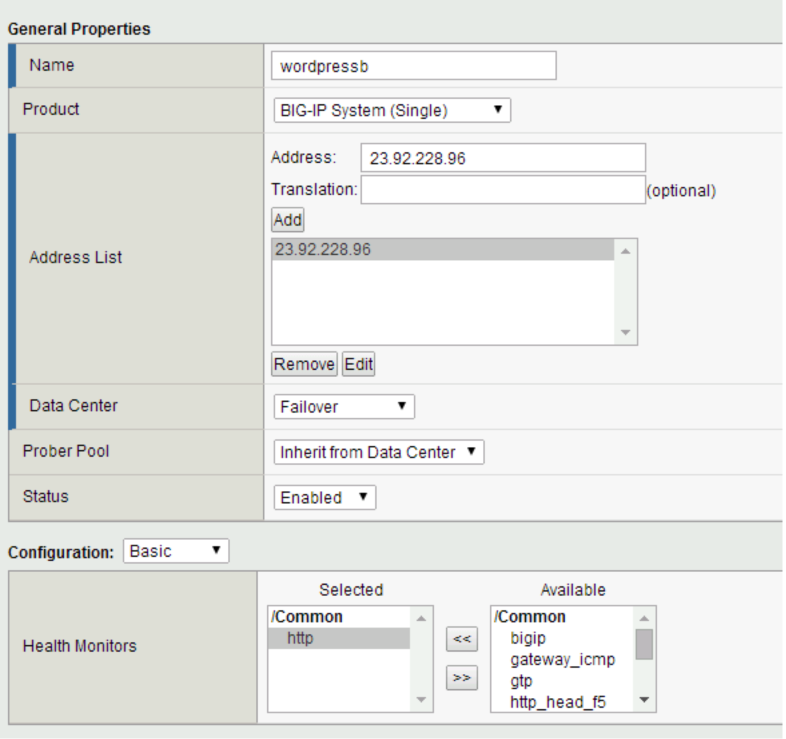

General Properties

Name	wordpressb
Product	BIG-IP System (Single)
Address List	Address: 23.92.228.96 Translation: (optional) 23.92.228.96
Data Center	Failover
Prober Pool	Inherit from Data Center
Status	Enabled

Configuration: Basic

Health Monitors — Selected: /Common http; Available: /Common bigip, gateway_icmp, gtp, http_head_f5

Figure 8 Failover WordPress Virtual Server Configuration

5. Create GTM Pools

From the **DNS** menu navigate to **GSLB -> Pools** Click ‘**Create’** to configure the load balancing pool for the WordPress servers. Your failover WordPress server will be in an inactive state it will therefore select the primary site until a failover occurs and the secondary server is activated. BIG-IP GTM will always direct traffic to the available server based on the load balancing algorithm selected (Global Availability). After naming the pool add the WordPress A and WordPressB servers to the pool. When you click on the ‘Manage’ button on the pools your screen should look like Figure 9:

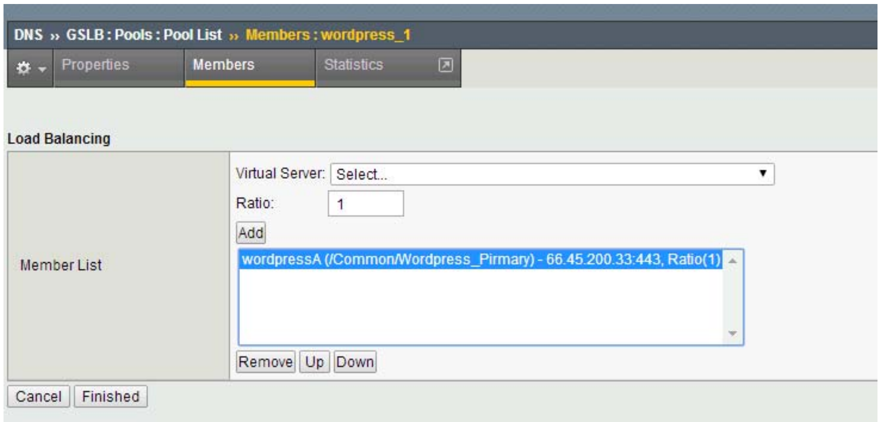

DNS » GSLB : Pools : Pool List » Members : wordpress_1

Load Balancing

Member List

Virtual Server: Select...

Ratio: 1

wordpressA (/Common/Wordpress_Pirmary) - 66.45.200.33:443, Ratio(1)

Figure 9: Add WordPress Virtual Servers to the Pool

6. Create DNS Listeners

In order for GTM to accept DNS requests you will need to create two listeners on port 53 for accepting DNS requests on the GTM. From the **DNS** tab navigate to **Delivery -> Listeners.** Create both a TCP and UDP listener using a Self IP address on the GTM. In this instance shown in Figure 10, we used the public Self IP address on the BIG-IP.

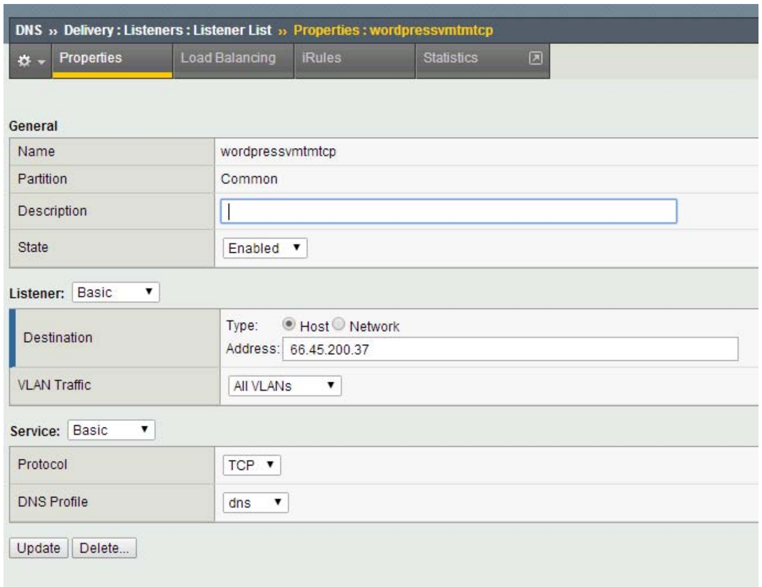
DNS » Delivery : Listeners : Listener List » Properties : wordpressvmtmtcp

General

Name	wordpressvmtmtcp
Partition	Common
Description	
State	Enabled

Listener: Basic

Destination	Type: ● Host ○ Network Address: 66.45.200.37
VLAN Traffic	All VLANs

Service: Basic

Protocol	TCP
DNS Profile	dns

Figure 10: Create DNS Listeners

7. Create Wide-IP

After creating the necessary server objects the wide-IP must be configured with the FQDN that you will be using for DNS requests. There are many configuration options within F5 GTM Wide-IPs that are outside the scope of this document. For additional information on GTM Wide-IP configuration settings please refer to the setup guides for GTM on https://support.f5.com.

To create the Wide-IP for the WordPress servers open the **DNS** menu and navigate to **GSLB -> Wide IP** click 'Create' to configure a new Wide IP. For the WordPress GSLB configuration we set the following parameters in the table below:

Name	.xxx.xxxx (the FQDN of WordPress)
Description	Optional (WordPress Deployment)
iRules	Leave Blank
Load Balancing Method (Pools Section)	Select Global Availability
Persistence	Disabled
Pool	'Add' Wordpress_1
Last Resort Pool	None

Click 'Finished' committing the changes.

8. Add BIGIP02 to Sync Group

As we have deployed the BIGIP01 device in the primary on-premises location, we configure BIGIP02 as a secondary GTM device within a vCloud Hybrid Service location to maintain redundancy and availability in the event that a Failover is required. This section requires that you have added the secondary GTM BIGIP02 device to your sync group by running the gtm_add script from the command line of the BIG-IP. This will exchange SSL key information between the two devices and will allow configuration data to be synchronized. You will also need the GTM server object to be created on the primary GTM.

As shown in Figure 10, from **DNS** menu, navigate to **Settings -> GSLB -> General** check 'Synchronize' and designate the group name for your synchronization. Click 'Update'. Verify synchronization changes by logging in to the secondary BIIP02 device; Wide-IPs, pools and other objects should be available on the secondary device.

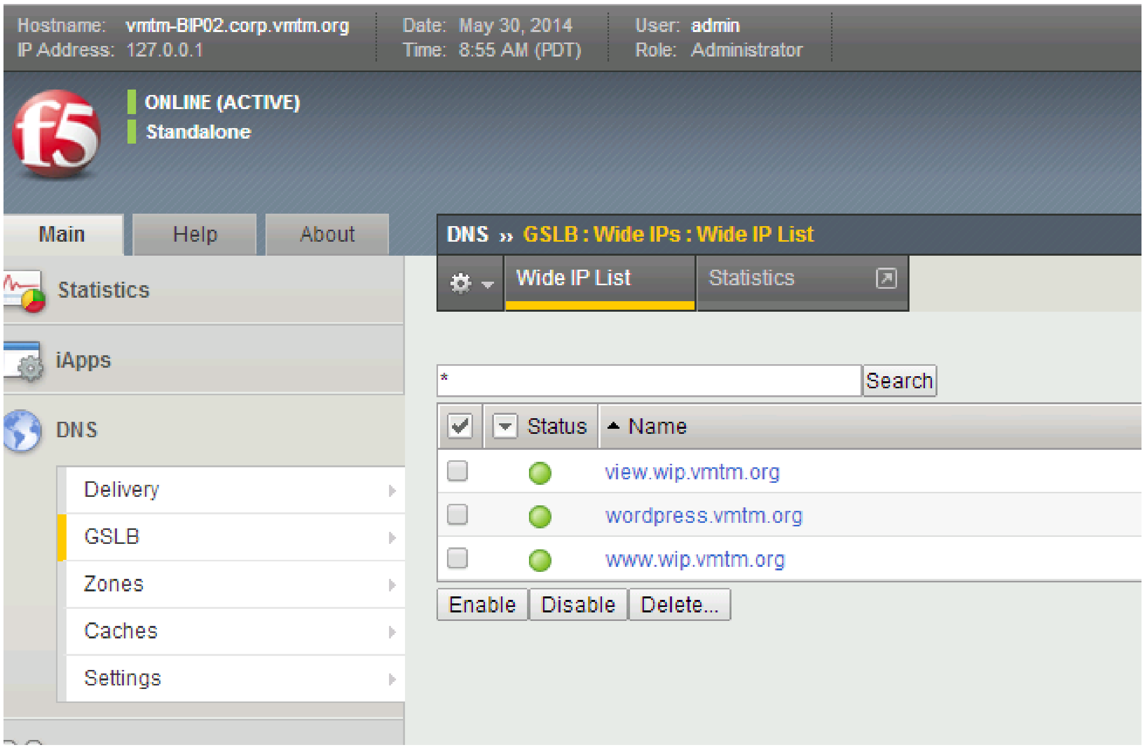
Hostname: vmtm-BIP02.corp.vmtm.org
IP Address: 127.0.0.1
Date: May 30, 2014
Time: 8:55 AM (PDT)
User: admin
Role: Administrator

ONLINE (ACTIVE)
Standalone

DNS » GSLB : Wide IPs : Wide IP List

☑	Status	Name
☐	●	view.wip.vmtm.org
☐	●	wordpress.vmtm.org
☐	●	www.wip.vmtm.org

Figure 11: Wide-IPs on BIP02

9. Delegate DNS Traffic to Wide-IPs on BIG-IP GTM

BIG-IP (GTM™) resolves DNS queries that match a Wide IP name. BIG-IP GTM can work in conjunction with an existing DNS server on your network. In this situation, you configure the DNS server to delegate Wide IP-related requests to BIG-IP GTM for name resolution.

In this implementation we have already configured and setup the Wide-IP on the BIGIP01 GTM device and synchronized changes to the BIGIP02 secondary GTM device. The local DNS server will be configured to allow all queries for the above Wide-IP to be redirected to the GTM. The following configurations were made on the DNS server in the environment as shown in Figure 12:

1. Create an A record for each GTM that defines the domain name for each GTM to in the sync group. In the above example we create (A) records for both BIP01.companyname.com and BIP02.companyname.com.
2. Create a Name Server Record (NS) that defines the delegated zone for which the GTMs are responsible. In this example we created a subzone **wip.companyname.com** and used the above records for the name servers to serve this sub zone.
3. Create a CNAME on your DNS server that forwards requests for **www.companyname.com** to a managed record called www.wip.companyname.com (This record is owned and served by the GTM's).

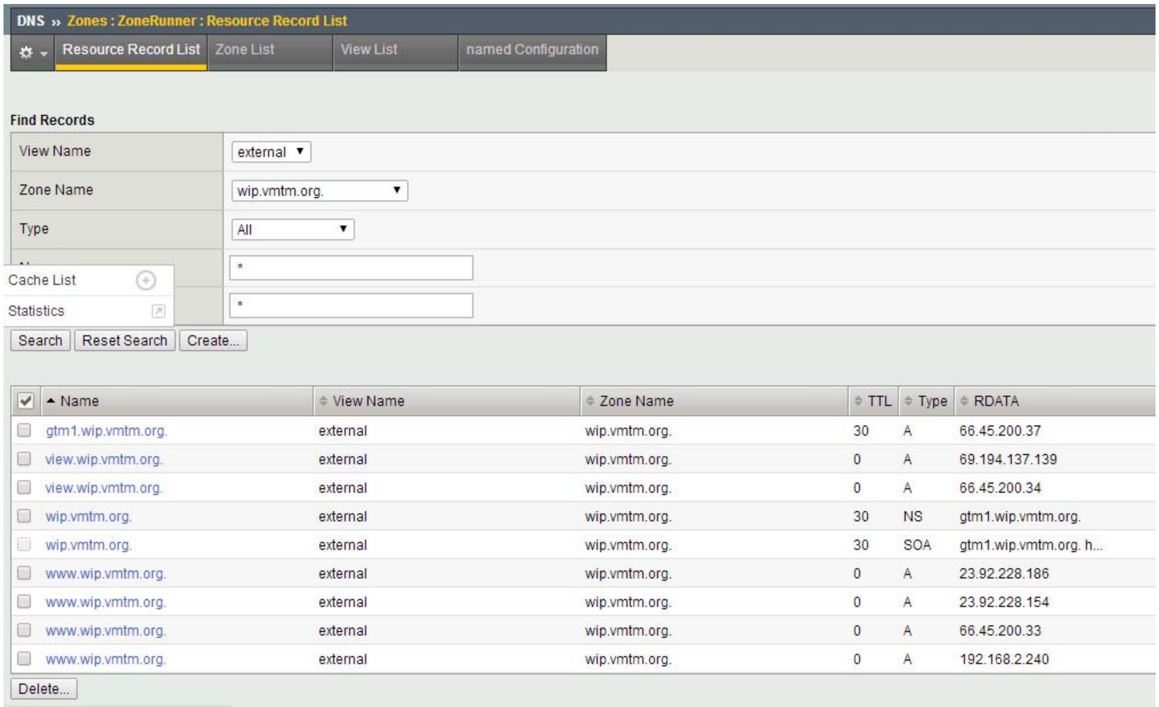

Find Records

View Name	external
Zone Name	wip.vmtm.org.
Type	All
	*
	*

☑	Name	View Name	Zone Name	TTL	Type	RDATA
□	gtm1.wip.vmtm.org.	external	wip.vmtm.org.	30	A	66.45.200.37
□	view.wip.vmtm.org.	external	wip.vmtm.org.	0	A	69.194.137.139
□	view.wip.vmtm.org.	external	wip.vmtm.org.	0	A	66.45.200.34
□	wip.vmtm.org.	external	wip.vmtm.org.	30	NS	gtm1.wip.vmtm.org.
□	wip.vmtm.org.	external	wip.vmtm.org.	30	SOA	gtm1.wip.vmtm.org. h...
□	www.wip.vmtm.org.	external	wip.vmtm.org.	0	A	23.92.228.186
□	www.wip.vmtm.org.	external	wip.vmtm.org.	0	A	23.92.228.154
□	www.wip.vmtm.org.	external	wip.vmtm.org.	0	A	66.45.200.33
□	www.wip.vmtm.org.	external	wip.vmtm.org.	0	A	192.168.2.240

Figure 12: DNS records from ZoneRunner interface on BIG-IP

C. Validating Configuration

After completing the above steps the validation of the primary site via the FQDN took place. Open a browser and point it to https:// xxx.xxx.xxx you will be directed to a basic WordPress web page. Open a command line and perform an nslookup xxx.xxx.xxx it will reflect the primary server is active and return the public IP address of the primary onsite location.

1. Execute Failover to vCloud Hybrid Service – Disaster Recovery

For the purposes of this we used the "Test" failover feature within the service, but the full failover would be the same.

- Click the 'Test Failover' button in the vSphere Web Client

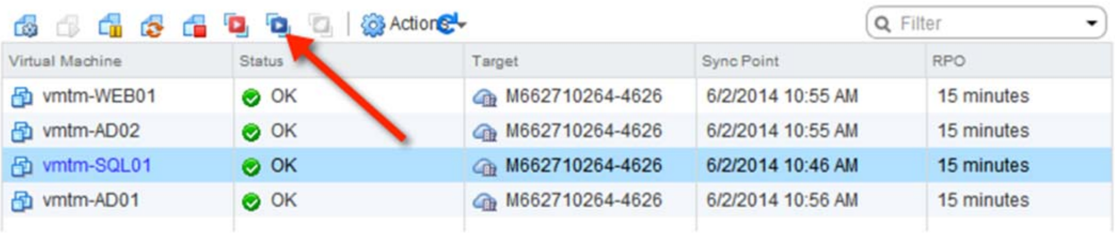

Virtual Machine	Status	Target	Sync Point	RPO
vmtm-WEB01	OK	M662710264-4626	6/2/2014 10:55 AM	15 minutes
vmtm-AD02	OK	M662710264-4626	6/2/2014 10:55 AM	15 minutes
vmtm-SQL01	OK	M662710264-4626	6/2/2014 10:46 AM	15 minutes
vmtm-AD01	OK	M662710264-4626	6/2/2014 10:56 AM	15 minutes

- Assign a new internal IP. This is usually handled by the static IP pool, DHCP, or manual assignment. The machine will be automatically connected to the network configured as the “Test Network” but can be manually edited.

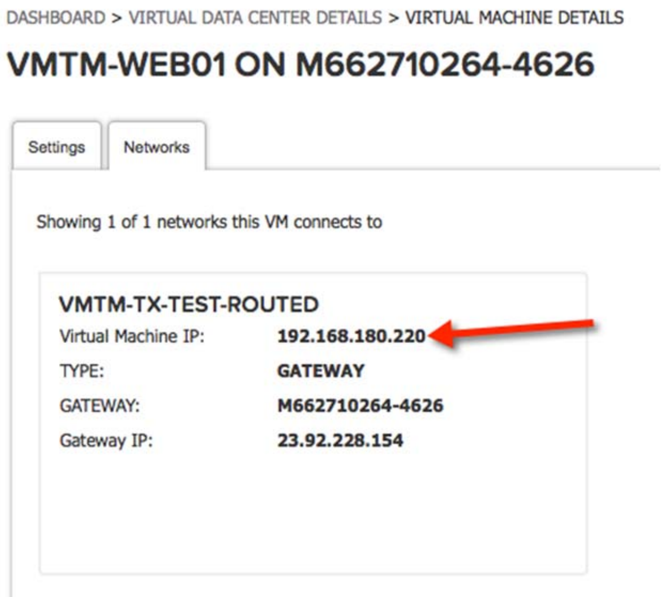

DASHBOARD > VIRTUAL DATA CENTER DETAILS > VIRTUAL MACHINE DETAILS

VMTM-WEB01 ON M662710264-4626

Showing 1 of 1 networks this VM connects to

VMTM-TX-TEST-ROUTED

Virtual Machine IP:	192.168.180.220
TYPE:	GATEWAY
GATEWAY:	M662710264-4626
Gateway IP:	23.92.228.154

- Create DNAT and firewall rules for the public IP address to pass traffic to the virtual machine’s new private IP. It is important to note that the firewall must also be configured to allow outbound traffic and include a SNAT rule for the subnet or machine that is being failover over.

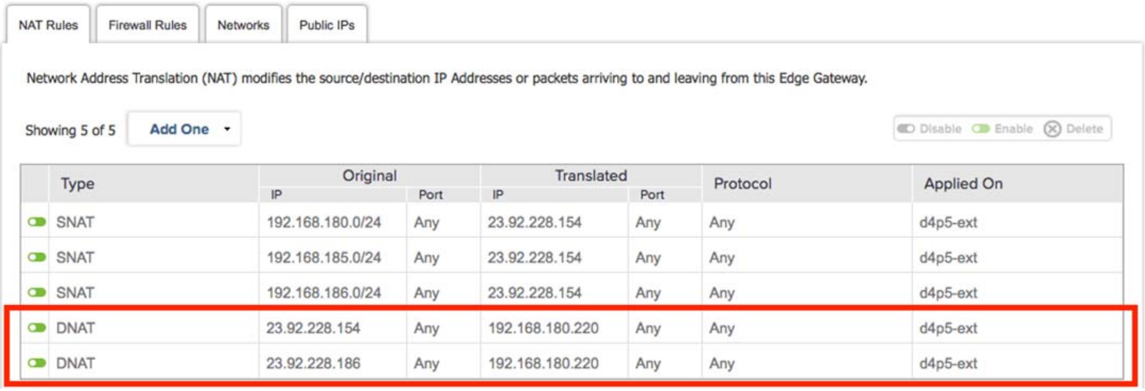

Network Address Translation (NAT) modifies the source/destination IP Addresses or packets arriving to and leaving from this Edge Gateway.

Showing 5 of 5

Type	Original		Translated		Protocol	Applied On
	IP	Port	IP	Port		
SNAT	192.168.180.0/24	Any	23.92.228.154	Any	Any	d4p5-ext
SNAT	192.168.185.0/24	Any	23.92.228.154	Any	Any	d4p5-ext
SNAT	192.168.186.0/24	Any	23.92.228.154	Any	Any	d4p5-ext
DNAT	23.92.228.154	Any	192.168.180.220	Any	Any	d4p5-ext
DNAT	23.92.228.186	Any	192.168.180.220	Any	Any	d4p5-ext

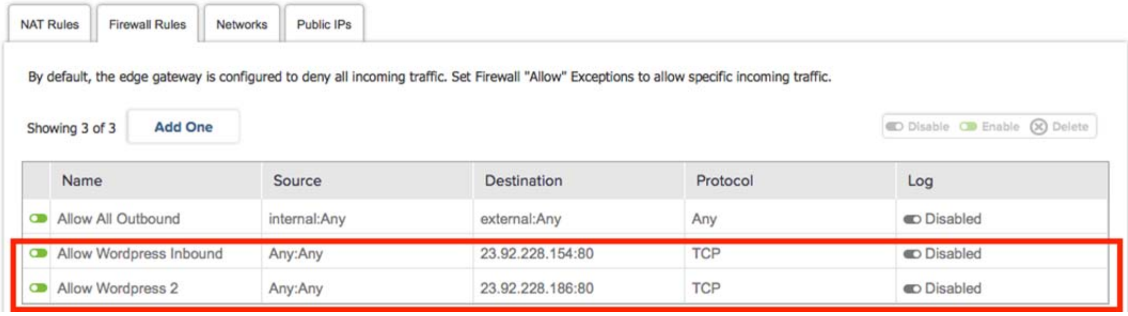

By default, the edge gateway is configured to deny all incoming traffic. Set Firewall "Allow" Exceptions to allow specific incoming traffic.

Name	Source	Destination	Protocol	Log
Allow All Outbound	internal:Any	external:Any	Any	Disabled
Allow Wordpress Inbound	Any:Any	23.92.228.154:80	TCP	Disabled
Allow Wordpress 2	Any:Any	23.92.228.186:80	TCP	Disabled

Figure 13: Create firewall rules

2. Power on the Virtual Machine

Generally you can pre-create the rules and adjust the internal IP settings during the test failover and re-use them in the event of a full failover. This way the IP's may change in the rules but the rules themselves are still in place.

After the failover has taken place successfully and the primary site IP address is offline open a browser and navigate to http://www.companyname.com. The same landing page will be visible. Open a command line and run nslookup www.companyname.com it should reflect now that the public IP of the failover site is now servicing requests.

This demo site is used by the VMware vCloud Hybrid Service Technical Marketing team for videos, documentation and general demonstrations of various hybrid cloud concept.s.

Figure 14: WordPress Demo Site on Primary

D. Create Multi-Site Availability of View Connection Servers via BIG-IP GTM

In the second example, we show global availability of multiple Horizon View Connection servers via BIG-IP GTM. Instead of Global Availability, we selected Round Robin as the LB algorithm to demonstrate availability across two separate vCloud Hybrid Service instances. Repeat the above steps with the following differences:

1. Create server objects and virtual server objects for the view connection servers define these server objects as with the same parameters as in the above example.

2. Create a global load balancing pool and add both virtual servers to the pool. In the example we created a pool call '*Horizon_Availability'*.
3. Create a Wide-IP in this example we have named the Wide-IP view.wip.companyname.com. Add the new Horizon_Availability pool to the configuration. Under load balancing algorithm choose 'Round Robin' this will demonstrate connections to the primary and secondary pool members.
4. In DNS we created a CNAME alias for view.companyname.com that points to view.wip.companyname.com
5. Open a browser and connect to https://view.companyname.com
6. Open a command line and run nslookup view.companyname.com
7. Close your browser and rebrowse to https://view.companyname.com
8. Open a command line and run nslookup view.companyname.com
9. Your address will read the secondary IP address.

E. Conclusion:

In the above scenarios, we demonstrated utilizing BIG-IP Global Traffic Manager to provide DNS services in support of a vCloud Hybrid Service DR and global application availability. It is important to note that any application which requires DNS services in vCloud Hybrid Service could benefit from utilizing BIG-IP GTM to ensure application availability, disaster recovery and a high performing user experience. Federating application resources between main data center environments and the cloud can provide enterprise customers the flexibility and efficiency that public cloud provides while maintaining a single pane of glass management capability.

Learn More:

F5 VMware Partner Page: https://f5.com/partners/product-technology-alliances/vmware
To Request a Free Trial of BIG-IP GTM, send an email to vchstrial@f5.com
See more details on F5 on vCloud Hybrid Service marketplace [link]

F5 Networks, Inc. 401 Elliott Avenue West Seattle, WA 98119. USA, Tel 888-88BIGIP. FAX 206-272-5556. www.f5.com

VMware Inc, 3401 Hillview Avenue, Palo Alto, CA 94304. USA, Tel 877-486-9273. FAX 650-427-5001. www.vmware.com