

Integrating IBM® Multi-site Workload Lifeline with F5 BIG-IP®

Grant Mericle, IBM
Divya Gundamaraju, F5 Networks
August 2018

Contents

- Introduction 1
- Technology Brief 1
 - F5 BIG-IP Local Traffic Manager 1
 - IBM Multi-site Workload Lifeline 2
- Use Cases 4
 - Use Case 1: Proof of Concept..... 4
 - Use Case 2: F5 redundancy and client access via hostname or IP address 14

Introduction

z/OS systems have been used for decades to provide a robust operating environment for banking, insurance, and general commercial computing needs. Regulatory and business requirements have driven the development of extensive disaster recovery solutions across multiple physical sites, often located hundreds or even thousands of kilometers apart. Failover to a backup site can take multiple hours before business applications are available for use. While a multi-hour outage may be acceptable for some applications, there are others for which a customer needs continuous availability. IBM Multi-site Workload Lifeline works with the F5 BIG-IP Local Traffic Manager to make Continuous Availability a reality for these critical business workloads.

Technology Brief

F5 BIG-IP Local Traffic Manager

Applications drive innovation and profitability, allowing businesses to leverage trends such as cloud computing, mobility, and software-defined networks (SDNs). BIG-IP® Local Traffic Manager™ (LTM) is F5’s enterprise-class layer 4-7 intelligent load balancer that can optimally manage network traffic so applications are always fast, available, and secure. With its scalable and programmable infrastructure, BIG-IP LTM provides extensibility and flexibility of application services that are needed to manage physical, virtual, and cloud infrastructure.

One of the core functionalities of BIG-IP is application health monitoring. BIG-IP LTM monitors determine the availability and performance of devices, links, and services on a network. If a monitored device, link, or service does not respond within a specified timeout period, or returns a status indicating that performance is degraded or that the load is excessive, the BIG-IP® system can redirect the traffic to another resource. BIG-IP LTM's Server/Application State Protocol (SASP) monitor type is an implementation of SASP that provides a mechanism for BIG-IP and workload management systems to communicate better ways of distributing the existing workload to the group members. In this context, SASP monitor communicates with a Group Workload Manager (GWM) to query for information on the current weights of each managed resource. These weights determine which resource currently provides the best response time. When the monitor receives this information from the GWM, it configures the dynamic ratio option for the resources, allowing the BIG-IP system to select the most appropriate resource to respond to a connection request.

IBM Multi-site Workload Lifeline

IBM Multi-site Workload Lifeline consists of two components: Lifeline Advisors and Lifeline Agents. The purpose of the product is to monitor two IBM Z data centers and seamlessly reroute workload connections when the server or site fails. A Lifeline Agent runs on each z/OS system on both sites. One Lifeline Advisor, configured as a primary or backup, runs in each site. The Lifeline Agents send server availability and health information to the primary Lifeline Advisor. The primary Lifeline Advisor provides systems administrators with a central location for determining server and state status, and controlling the routing of each configured workload.

Combined Solution

With a Lifeline & BIG-IP solution, clients connect to a BIG-IP vserver instead of directly to a datacenter server. The z/OS administrator issues commands to the primary Lifeline Advisor to select the active site for a workload. The Lifeline Advisor acts as the GWM, communicating with the BIG-IP LTM via the SASP protocol to instruct the BIG-IP to route all workload connections to the selected site. When either a planned or unplanned site outage occurs, the Lifeline Advisor can switch the workload to the alternate site in minutes. Figure 1 shows a high-level overview of a Lifeline and BIG-IP environment.

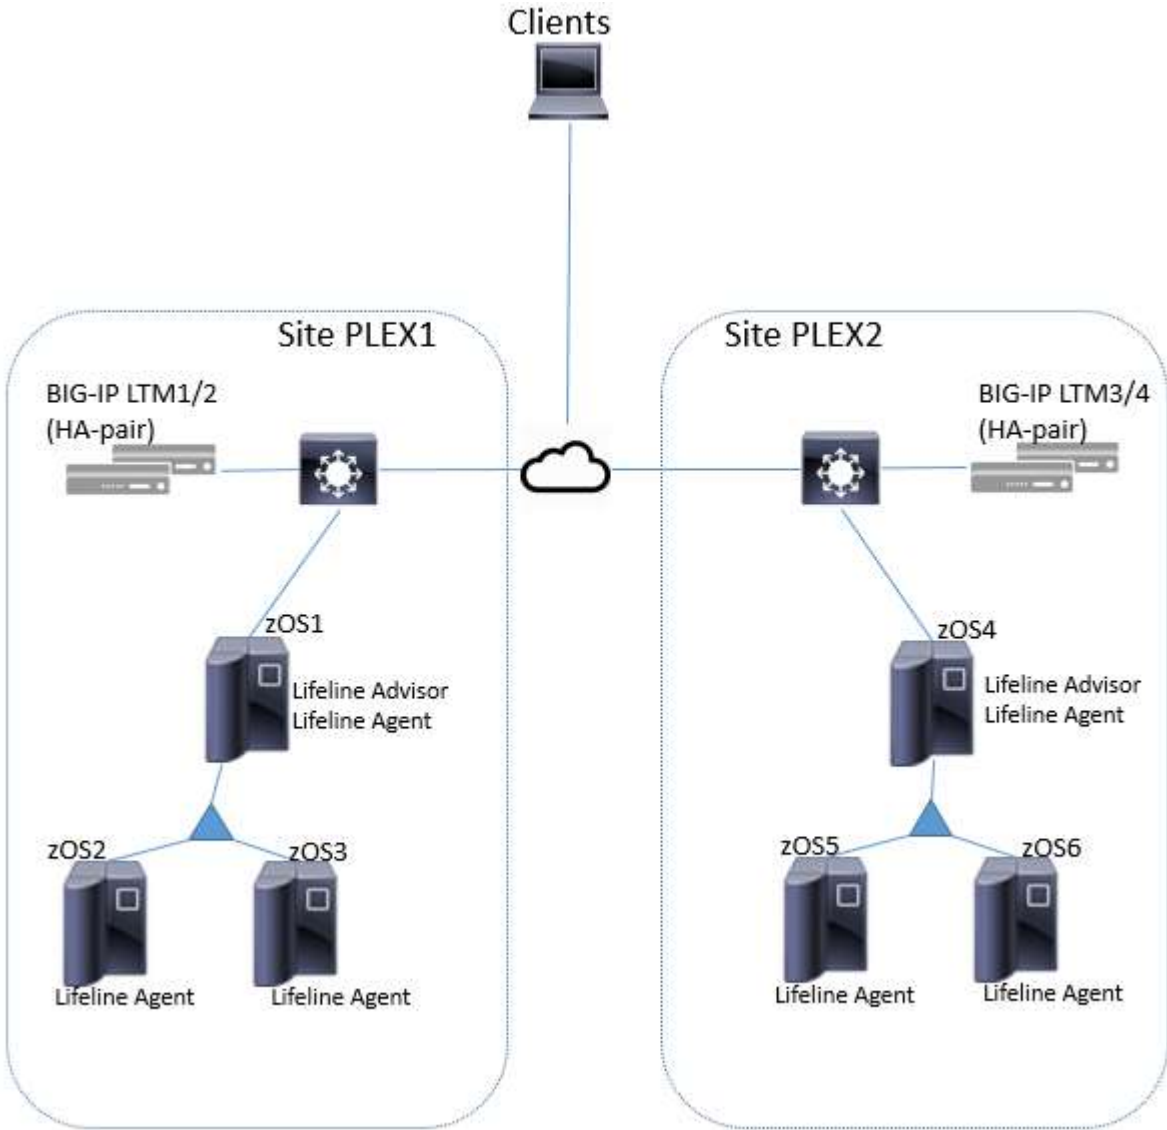


Figure 1: High-Level Architecture of BIG-IP and Multi-site Workload Lifeline

Assumptions

In the world of Enterprise Computing, IBM Z administration and Network (e.g. routers, switches) administration are typically handled by distinct groups. This presents a challenge for customers that are implementing a solution that crosses ownership boundaries. This document provides an architectural overview and real-world examples that touch on both the z/OS as well as the networking components. The target audience is IBM Z architects and Network architects.

Prerequisites

Infrastructure

Each use case includes more specific infrastructure requirements, but these general requirements apply to all environments:

- **2 z/OS sysplexes**
- **F5 BIG-IP LTM**
- **Layer 3 Network connectivity between all systems**

Use Cases

The first use case introduces the solution components. It is intended to be a minimum viable installation for proof-of-concept use. The other use cases cover various common configurations. The intention is to find a use case that maps most closely to your workload and gain insight on configuration details and best practices.

Use Case 1: Proof of Concept

This use case is not intended to be a model for production environments, but rather to provide a testbed for exploring the interaction between Lifeline components and the BIG-IP. A single telnet workload will be defined to allow client connections to be run through the BIG-IP to the active z/OS site.

Infrastructure

To deploy this solution, the following infrastructure prerequisites are required:

- Client system with telnet client (any platform)
- Two z/OS systems or sysplexes
- IBM Multi-site Workload Lifeline v2.5
- BIG-IP LTM v13.0 physical or virtual edition instance in a one-armed configuration
 - In a fully redundant solution, BIG-IP would be deployed in an HA-pair configuration
- IP connectivity (layer 3) between all z/OS systems and BIG-IP

Lab Environment:

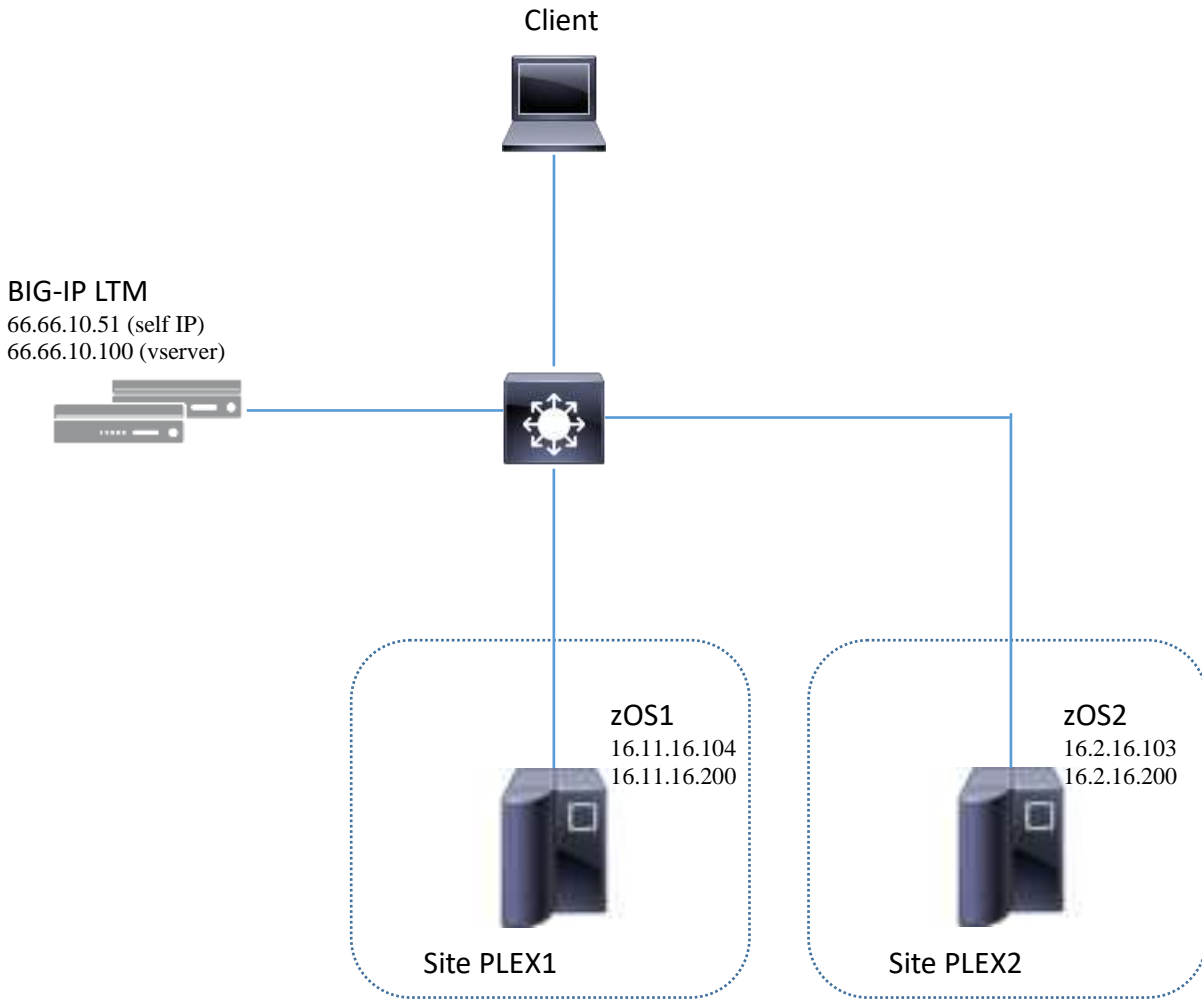


Figure 2: Simplified site architecture for use case 1

z/OS TCP/IP Configuration

In this use case, each z/OS system represents a separate sysplex. In a real-world deployment, each sysplex would likely be located in its own data center. Each sysplex must run one instance of Lifeline Advisor for redundancy purposes. Each z/OS system must run an instance of Lifeline Agent to monitor server availability and health for the workload. In this minimal 2-system deployment, the Lifeline Advisor and Agent will run on the same system.

On each z/OS system, define a static VIPA. The static VIPA is used to when creating Lifeline Agent to Lifeline Advisor connections.(Use full clause to increase the clarity) In each z/OS sysplex, define a Dynamic VIPA that distributes port 23. This Dynamic VIPA will be used for workload connections.

zOS1 (sysplex PLEX1)

Static VIPA 16.11.16.104

Dynamic VIPA 16.11.16.200

VIPADISTRIBUTE 16.11.16.200 PORT 23 DESTIP ALL

zOS2 (sysplex PLEX2)

Static VIPA 16.2.16.103

Dynamic VIPA 16.2.16.200

VIPADISTRIBUTE 16.2.16.200 PORT 23 DESTIP ALL

*Note: In this use case the static and dynamic VIPAs on each system are created in the same subnet for simplicity; this is not a hard requirement.

Verification:

```
zOS1:/u/home/gsm>netstat -B 16.11.16.200+23 -O
MVS TCP/IP NETSTAT CS V2R2          TCPIP Name: TCPSVT          14:00:04
Dynamic VIPA Destination Port Table for TCP/IP stacks:
Dest:      16.11.16.200..23
  DestXCF:  199.11.80.104
  TotalConn: 0000000003 Rdy: 001 WLM: 14 TSR: 100
  DistMethod: ServerWLM
  Flg:
```

```
zOS2:/u/home/gsm>netstat -B 16.2.16.200+23 -O
MVS TCP/IP NETSTAT CS V2R3          TCPIP Name: TCPSVT          14:02:24
Dynamic VIPA Destination Port Table for TCP/IP stacks:
Dest:      16.2.16.200..23
  DestXCF:  199.2.80.103
  TotalConn: 0000000000 Rdy: 001 WLM: 00 TSR: 100
  DistMethod: ServerWLM
  Flg:
```

At this point, test connections from the client system to either DVIPA should succeed. Generally, "At this point" is needless and doesn't add meaning to the text.

Configure the Primary Lifeline Advisor

Use the sample procedure included with the installed Lifeline product, then set the Advisor configuration as follows:

```
debug_level          127
advisor_takeover_policy  automatic
advisor_timeout      10
workload_switch_policy manual
failure_detection_interval 30
update_interval      10

#Primary_GWM
lb_connection_v4     16.11.16.104..3860
```

```

#Secondary_GWM
peer_advisor_id    16.2.16.103

lb_id_list
{
  66.66.10.51 #Big-IP self IP
}

agent_connection_port    8181
agent_id_list
{
  16.11.16.104..4000
  16.2.16.103..4000
  #Include the static VIPA of each z/OS system where a Lifeline Agent
  #is running
}

advisor_connection_port 8000
advisor_id_list
{
  16.11.16.104
  16.2.16.103
}

cross_sysplex_list
{
  #WORK1 definition consists of the DVIPA on each site:
  16.11.16.200..23,PLEX1,WORK1
  16.2.16.200..23,PLEX2,WORK1
}

```

You can find more z/OS configuration steps (i.e. RACF, etc) to authorize each Lifeline process to start in the IBM Multi-site Workload Lifeline manual (<https://www.ibm.com/software/network/lifeline/library/>).

Start the Lifeline Advisor in primary mode and issue MODIFY AQSADV,DISPLAY,ADVISOR on that system to verify the configuration:

```

AQS0141I ADVISOR SUMMARY
ADVISOR ROLE   : PRIMARY
IPADDR        : 16.11.16.104
LOAD BALANCERS:
  IPADDR      : 66.66.10.51           CONNECTED: NO
  TIER        : UNKNOWN
AGENTS        :
  IPADDR      : 16.11.16.104         CONNECTED: NO
  IPADDR      : 16.2.16.103         CONNECTED: NO
PEER ADVISOR   :
IPADDR        : NONE

```

The Secondary Lifeline Advisor configuration is identical to the primary, with this exception:

```

#Primary_GWM
lb_connection_v4    16.2.16.103..3860
#Secondary_GWM
peer_advisor_id    16.11.16.104

```

Start the Lifeline Advisor on zOS2 in secondary mode. Verify that it connects to the primary Lifeline Advisor by looking for these messages in the system or job log:

```

AQS0101I LLADVSR STARTING
AQS0173I ADVISOR IS NOW IN SECONDARY ROLE
AQS0129I LLADVSR CONNECTED TO ADVISOR AT 16.11.16.104

```

Configure a Lifeline Agent on zOS1 and zOS2.

zOS1 Lifeline Agent

```
debug_level 127

advisor_id_list
{
  16.11.16.104..8181
  16.2.16.103..8181
}

# host_connection MUST match an entry in the
# agent_id_list statement in the advisor configuration file
host_connection 16.11.16.104..4000
site_name PLEX1
```

zOS2 Lifeline Agent

```
debug_level 127

advisor_id_list
{
  16.11.16.104..8181
  16.2.16.103..8181
}

# host_connection MUST match an entry in the
# agent_id_list statement in the advisor configuration file
host_connection 16.2.16.103..4000
site_name PLEX2
```

Start both Lifeline Agents and reissue the MODIFY AQSADV,DISPLAY,ADVISOR command on the primary Lifeline Advisor system to verify connectivity:

```
AQS0141I ADVISOR SUMMARY
ADVISOR ROLE : PRIMARY
IPADDR      : 16.11.16.104
LOAD BALANCERS:
  IPADDR : 66.66.10.51          CONNECTED: NO
  TIER   : UNKNOWN
AGENTS      :
  IPADDR : 16.11.16.104        CONNECTED: YES
  IPADDR : 16.2.16.103        CONNECTED: YES
PEER ADVISOR :
IPADDR      : 16.2.16.103
```

BIG-IP configuration

Overview

To handle incoming client connections and route them to the correct z/OS site,

1. create a SASP monitor that will communicate with the Lifeline Advisor
2. create a pool that contains the Dynamic VIPA on each site
3. associate the SASP monitor with the pool
4. create a vserver listener and specify the pool created in step 2

Create a SASP monitor and point it to the Primary and Secondary Lifeline Advisors

Local Traffic » Monitors » SASP-F5

⚙️ Properties Instances

General Properties

Name	SASP-F5
Partition / Path	Common
Description	
Type	SASP
Parent Monitor	sasp

Configuration

Mode	Push ▼
GWM Primary Address	16.11.16.104
GWM Secondary Address	16.2.16.103
GWM Service Port	3860 Other: ▼
GWM Protocol	TCP ▼

Create a server pool consisting of the Dynamic VIPA addresses on each site. Configure it to use the SASP monitor created in the previous step. Specify a Load Balancing Method of 'Dynamic Ratio (member)'. For any additional pools created for this or other Lifeline-enabled workloads, you can reuse the same SASP monitor.

Local Traffic » Pools : Pool List » **New Pool...**

Configuration: Basic

Name	WORK1-POOL
Description	
Health Monitors	<div style="display: flex; justify-content: space-around;"> <div style="border: 1px solid #ccc; padding: 5px;"> <p>Active</p> <ul style="list-style-type: none"> <i>/Common</i> <li style="background-color: #007bff; color: white;">SASP-F5 </div> <div style="border: 1px solid #ccc; padding: 5px;"> <p>Available</p> <ul style="list-style-type: none"> <i>/Common</i> LL-SASP gateway_icmp http http_head_f5 </div> </div>

Resources

Load Balancing Method	Dynamic Ratio (member)
Priority Group Activation	Disabled
New Members	<p> <input checked="" type="radio"/> New Node <input type="radio"/> New FQDN Node <input type="radio"/> Node List </p> <p>Node Name: _____ (Optional)</p> <p>Address: 16.2.16.200</p> <p>Service Port: 23 TELNET</p> <p><input type="button" value="Add"/></p> <div style="border: 1px solid #ccc; padding: 5px; min-height: 40px;"> <p>R:1 P:0 C:0 16.11.16.200 16.11.16.200 :23</p> <p>R:1 P:0 C:0 16.2.16.200 16.2.16.200 :23</p> </div> <p><input type="button" value="Edit"/> <input type="button" value="Delete"/></p>

Create a vservers that uses the pool.

Source Address Translation should be set to Auto-Map.

The Destination Address is the IP address to be targeted by all clients accessing this workload service.

Local Traffic » Virtual Servers : Virtual Server List » WORK1-VS

Properties Resources Statistics

General Properties

Name	WORK1-VS
Partition / Path	Common
Description	
Type	Standard
Source Address	0.0.0.0/0
Destination Address/Mask	66.66.10.100
Service Port	23 TELNET
Notify Status to Virtual Address	<input checked="" type="checkbox"/>
Availability	Available (Enabled) - The virtual server is available
Syncookie Status	Off
State	Enabled

Source Address Translation: Auto Map

Default Pool: WORK1-POOL

Combined Environment Validation

As soon as the SASP Monitor is associated with a pool, BIG-IP will attempt to connect to the Primary Advisor. Once successful, zOS1 will display this message in the system and Lifeline Advisor job log:

AQS0133I LOAD BALANCER CONNECTED FROM 66.66.10.51

The Network Map on BIG-IP will show the vservers and pool in an Offline state:



Each pool address represents a different site. After the Lifeline Advisor starts, the workload is placed into a QUIESCED state. Confirm this initial state with a DISPLAY,WORKLOAD,DETAIL report from the primary Lifeline Advisor:

```
AQS0146I WORKLOAD DETAILS
TYPE : ACTIVE/STANDBY
WORKLOAD NAME : WORK1
STATE : QUIESCED
SITE : N/A
SERVERS:
IPADDR..PORT : 16.2.16.200..23
SYSTEM NAME : CHILE SITE : PLEX2 STATE : AVAIL
IPADDR..PORT : 16.11.16.200..23
SYSTEM NAME : RUSSIA SITE : PLEX1 STATE : AVAIL
1 OF 1 RECORDS DISPLAYED
```

In this QUIESCED state, all connection requests to the vserver will be reset. Recall that the Lifeline Advisor controls which site is considered 'Active' for a workload. Activate the workload to PLEX1 by issuing this command on the primary Lifeline Advisor system:

```
MODIFY AQSADV,ACTIVATE,WORKLOAD=WORK1,SITE=PLEX1
```

Lifeline Advisor responds with AQS0153I WORKLOAD WORK1 ACTIVATED ON SITE PLEX1

At this point, the BIG-IP Network Map shows the PLEX1 DVIPA (16.11.16.200:23) to be active:



The primary Lifeline Advisor shows similar information on the DISPLAY,WORKLOAD,DETAIL and DISPLAY,LB,DETAIL reports:

```
AQS0146I WORKLOAD DETAILS
TYPE : ACTIVE/STANDBY
WORKLOAD NAME : WORK1
STATE : ACTIVE
SITE : PLEX1
SERVERS:
IPADDR..PORT : 16.2.16.200..23
SYSTEM NAME : CHILE SITE : PLEX2 STATE : AVAIL
IPADDR..PORT : 16.11.16.200..23
SYSTEM NAME : RUSSIA SITE : PLEX1 STATE : AVAIL
1 OF 1 RECORDS DISPLAYED
```

```
AQS0113I LOAD BALANCER DETAILS
LB INDEX : 32 UUID : 34323261366563652D633662352D3133
62312D363032373538383563373739
```

```

IPADDR..PORT : 66.66.10.51..60232
HEALTH      : 7F          FLAGS      : NOCHANGE PUSH
GROUP NAME  : /COMMON/WORK1-POOL
WORKLOAD    : WORK1
GROUP FLAGS : CROSS_SYSPLEX
IPADDR..PORT: 16.2.16.200..23
SYSPLEX     : PLEX2
SYSTEM NAME: N/A          PROTOCOL  : TCP    AVAIL     : NO
WLM WEIGHT  : N/A          CS WEIGHT : N/A    NET WEIGHT: 00000
FLAGS       : CMQ DISTDVIPA
IPADDR..PORT: 16.11.16.200..23
SYSPLEX     : PLEX1
SYSTEM NAME: N/A          PROTOCOL  : TCP    AVAIL     : YES
WLM WEIGHT  : N/A          CS WEIGHT : N/A    NET WEIGHT: 00001
FLAGS       : DISTDVIPA
1 OF 1 RECORDS DISPLAYED

```

Use a telnet client to initiate connections to the vserver (66.66.10.100, port 23) and verify that each connection is routed to PLEX1.

/Common/LL-WORK100-POOL				Bits		Packets		Connections			
✓	▼ Status	▢ Pool	◆ Pool Member	◆ Partition / Path	◆ In	◆ Out	◆ In	◆ Out	◆ Current	◆ Maximum	◆ Total
<input type="checkbox"/>	●	LL-WORK100-POOL		Common	16.7K	26.1K	37	26	1	1	1
<input type="checkbox"/>	●		16.11.16.200:23	Common	16.7K	26.1K	37	26	1	1	1
<input type="checkbox"/>	◆		16.2.16.200:23	Common	0	0	0	0	0	0	0

Planned workload switch

To initiate a planned switch of the workload to the alternate site, first quiesce the workload by issuing this command on the primary Lifeline Advisor:

```
MODIFY AQSADV,QUIESCE,WORKLOAD=WORK1
```

- 12.32.28 f aqsadv,quiesce,work=work1
- 12.32.28 AQS0155I WORKLOAD WORK1 QUIESCED

At this point, the PLEX1 members are disabled and all new connection requests will fail. If the workload connections are short-lived, the operator can wait for all of them to complete before activating the workload to the PLEX2 site.

Before allowing the workload activation to complete on PLEX2, the Lifeline Advisor will check PLEX1 systems for any outstanding workload connections. If it finds any the activation will fail:

- 12.37.35 MODIFY AQSADV,ACTIVATE,WORKLOAD=WORK1,SITE=PLEX2
- 12.37.35 AQS0179I MODIFY ACTIVATE COMMAND IGNORED - WORKLOAD WORK1 HAS CONNECTIONS ON SITE PLEX1


The operator can either wait for the remaining connections to complete or forcibly terminate them using the deactivate command on the primary Lifeline Advisor:

- 12.39.33 MODIFY AQSADV,DEACTIVATE,WORKLOAD=WORK1
- 12.39.33 AQS0155I WORKLOAD WORK1 DEACTIVATED

The workload can now safely be activated to the PLEX2 site:

- 12.41.32 f aqsadv,act,work=work1,site=plex2
- 12.41.32 AQS0153I WORKLOAD WORK1 ACTIVATED ON SITE PLEX2

The PLEX2 members will be available on the BIG-IP and new connection requests will be routed to PLEX2.



Summary

In this simple example, you saw how BIG-IP uses the SASP protocol to set the status of its pool members. Lifeline provides a central point to monitor server availability on both sites and control which site is selected for all inbound connections for a defined workload.

Use Case 2: F5 redundancy and client access via hostname or IP address

This use case models a more typical production deployment. F5 redundancy is gained by using four BIG-IPs – two BIG-IPs in a device cluster deployed in each site. This redundancy presents some unique challenges, particularly for workloads where the client systems access a service via IP address rather than hostname. Two workloads will be defined, one accessed by clients via hostname and the other directly via IP address.

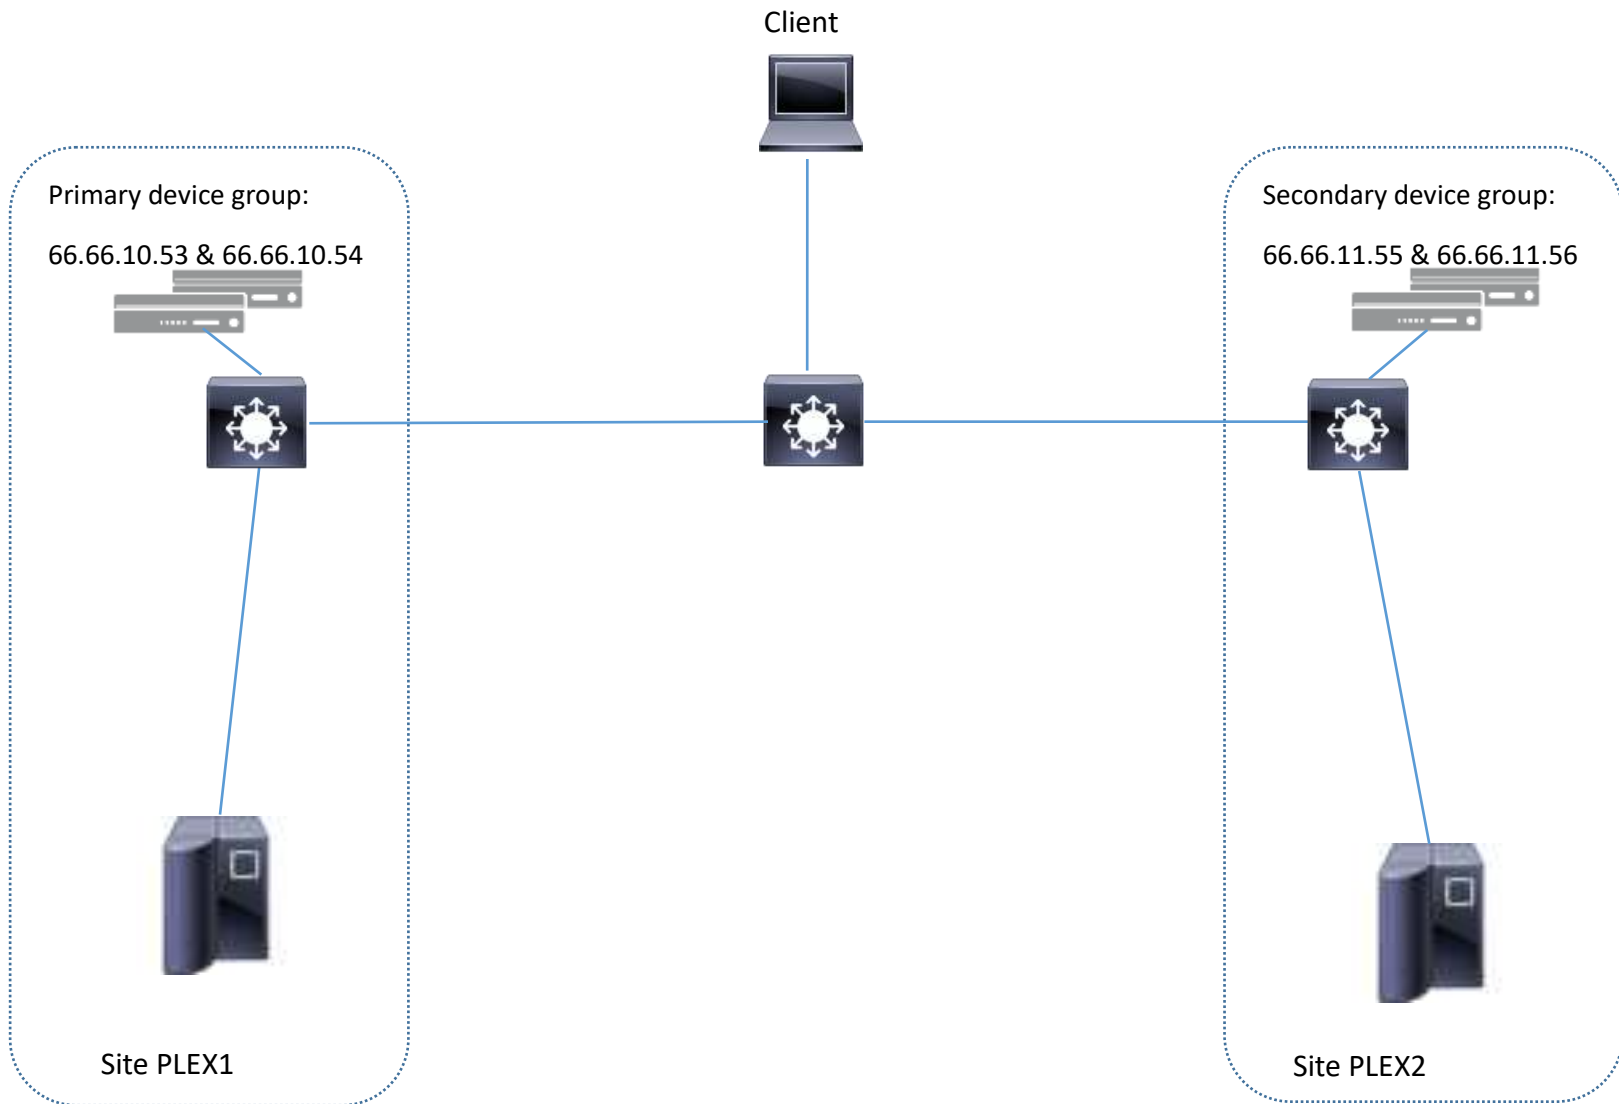
Infrastructure

This case uses the following infrastructure:

- Client system with telnet client and a web browser (any platform)
- Two z/OS systems or sysplexes, each running telnetd and a web server
- IBM Multi-site Workload Lifeline v2.5
- 4 BIG-IP LTM v13.0 physical or virtual edition instances

- IP connectivity (layer 3) between all z/OS systems and BIG-IPs

Lab Environment:



z/OS TCP/IP Configuration

The z/OS TCP/IP configuration is slightly expanded from use case 1 to support a second workload. Refer to the z/OS TCP/IP Configuration in Use Case 1. The only addition for use case 2 is to add port 80 to the VIPADISTRIBUTE definitions:

```
zOS1 (sysplex PLEX1)
    Dynamic VIPA 16.11.16.200
    VIPADISTRIBUTE 16.11.16.200 PORT 23 80 DESTIP ALL
```

```
zOS2 (sysplex PLEX2)
    Dynamic VIPA 16.2.16.200
```

VIPADISTRIBUTE 16.2.16.200 PORT 23 80 DESTIP ALL

Now configure a webserver on each z/OS system along with a default index page. It's helpful to customize the index page with the system or site name so that you can easily verify which site a connection was routed to.

Verification:

```
zOS1:/u/home/gsm>netstat -B 16.11.16.200+80 -o
MVS TCP/IP NETSTAT CS V2R2          TCPIP Name: TCPSVT
Dynamic VIPA Destination Port Table for TCP/IP stacks:
Dest:      16.11.16.200..80
  DestXCF:  199.11.80.104
  TotalConn: 0000000003 Rdy: 001  WLM: 14  TSR: 100
  DistMethod: ServerWLM
  Flg:
```

```
zOS2:/u/home/gsm>netstat -B 16.2.16.200+80 -o
MVS TCP/IP NETSTAT CS V2R3          TCPIP Name: TCPSVT
Dynamic VIPA Destination Port Table for TCP/IP stacks:
Dest:      16.2.16.200..80
  DestXCF:  199.2.80.103
  TotalConn: 0000000000 Rdy: 001  WLM: 00  TSR: 100
  DistMethod: ServerWLM
  Flg:
```

At this point, test connections from the client system web browser to either DVIPA should succeed.

Configure the Primary Lifeline Advisor

The Primary Lifeline Advisor configuration is identical to use case 1, except you'll define 4 BIG-IPs in the lb_id_list and second workload (WORK2) in the cross_sysplex_list definition.

```
lb_id_list
{
  66.66.10.53 #Big-IP self IP (HA cluster 1)
  66.66.10.54 #Big-IP self IP (HA cluster 1)
  66.66.11.55 #Big-IP self IP (HA cluster 2)
  66.66.11.56 #Big-IP self IP (HA cluster 2)
}

cross_sysplex_list
{
  #WORK1 definition consists of the DVIPA on each site:
  16.11.16.200..23, PLEX1, WORK1
  16.2.16.200..23, PLEX2, WORK1

  #WORK2 definition
  16.11.16.200..80, PLEX1, WORK2
  16.2.16.200..80, PLEX2, WORK2
}
```

Make the same additions to the Secondary Lifeline Advisor configuration.

Start the Lifeline Advisor in primary mode and issue MODIFY AQSADV,DISPLAY,ADVISOR on that system to verify the configuration:

```
AQS0141I ADVISOR SUMMARY
ADVISOR ROLE   : PRIMARY
IPADDR        : 16.11.16.104
LOAD BALANCERS:
  IPADDR      : 66.66.10.53          CONNECTED: NO
  TIER        : UNKNOWN
  IPADDR      : 66.66.10.54          CONNECTED: NO
  TIER        : UNKNOWN
  IPADDR      : 66.66.11.55         CONNECTED: NO
  TIER        : UNKNOWN
  IPADDR      : 66.66.11.56         CONNECTED: NO
  TIER        : UNKNOWN
AGENTS        :
  IPADDR      : 16.11.16.104        CONNECTED: NO
  IPADDR      : 16.2.16.103        CONNECTED: NO
PEER ADVISOR  :
IPADDR        : NONE
```

The Secondary Lifeline Advisor configuration is identical to the primary, with this exception:

```
#Primary_GWM
lb_connection_v4 16.2.16.103..3860
#Secondary_GWM
peer_advisor_id 16.11.16.104
```

Start the Lifeline Advisor on zOS2 in secondary mode. Verify that it connects to the primary Lifeline Advisor by looking for these messages in the system or job log:

```
AQS0101I LLADVSR STARTING
AQS0173I ADVISOR IS NOW IN SECONDARY ROLE
AQS0129I LLADVSR CONNECTED TO ADVISOR AT 16.11.16.104
```

Configure a Lifeline Agent on zOS1 and zOS2.

Lifeline Agent configurations are identical to use case 1.

Start both Lifeline Agents then issue the MODIFY AQSADV,DISPLAY,ADVISOR command on the primary Lifeline Advisor system to verify connectivity:

```
AQS0141I ADVISOR SUMMARY
ADVISOR ROLE   : PRIMARY
IPADDR        : 16.11.16.104
LOAD BALANCERS:
  IPADDR      : 66.66.10.53          CONNECTED: NO
  TIER        : UNKNOWN
  IPADDR      : 66.66.10.54          CONNECTED: NO
  TIER        : UNKNOWN
  IPADDR      : 66.66.11.55         CONNECTED: NO
  TIER        : UNKNOWN
  IPADDR      : 66.66.11.56         CONNECTED: NO
  TIER        : UNKNOWN
AGENTS        :
  IPADDR      : 16.11.16.104        CONNECTED: YES
  IPADDR      : 16.2.16.103        CONNECTED: YES
PEER ADVISOR  :
IPADDR        : 16.2.16.103
```

BIG-IP configuration

Redundancy

This use case has four BIG-IP devices. Two are deployed in each site to prevent workload failure if a whole site fails. Use the built-in configuration wizard to configure the device groups in each site. If you had layer-2 adjacency spanning both sites, you could configure all four BIG-IPs in a single device group.

Workload access by hostname

The telnet workload will be accessed via DNS, so you can use both device groups to service connections. You will create a unique vserver for each device group and associate both vservers with the DNS record. This could be accomplished in a standard DNS server or by using F5 DNS technology to obtain more intelligent load balancing between the device groups. The F5 DNS configuration is beyond the scope of this use case but guidance can be found on the F5 Knowledge Center.

LTM configuration

To handle incoming client connections and route them to the correct z/OS site,

1. create a SASP monitor that will communicate with the Lifeline Advisor
2. create a pool that contains the Dynamic VIPA on each site
3. associate the SASP monitor with the pool
4. create a vserver listener and specify the pool created in step 2

If you have already completed use case 1, the SASP monitor instance already exists. If not, create it now.

Create a SASP monitor and point it to the Primary and Secondary Lifeline Advisors

Local Traffic » Monitors » SASP-F5

⚙️ Properties Instances

General Properties

Name	SASP-F5
Partition / Path	Common
Description	
Type	SASP
Parent Monitor	sasp

Configuration

Mode	Push ▼
GWM Primary Address	16.11.16.104
GWM Secondary Address	16.2.16.103
GWM Service Port	3860 Other: ▼
GWM Protocol	TCP ▼

Follow the steps in use case 1 to create the pool vserver for the telnet (port 23) workload on the first device cluster (66.66.10.53 and 66.66.10.54). Use the virtual address 66.66.10.100 for this device cluster.

Repeat the steps to create the monitor, pool, and vserver for the telnet workload on the second HA cluster (66.66.11.55 and 66.66.11.56). Use a different virtual address (66.66.11.100) for this HA cluster.

Associate these vserver addresses with a hostname in DNS, as seen in the following example.

```
telnet.svt390.com    A      66.66.10.100
                   A      66.66.11.100
```

Combined Environment Validation

As soon as the SASP Monitor is associated with a pool, BIG-IP will attempt to connect to the Primary Advisor. When the connection is successful, zOS1 will display this message in the system and Lifeline Advisor job log:

```
AQS0133I LOAD BALANCER CONNECTED FROM 66.66.10.53
```

All four BIG-IPs should form connections with the Primary Lifeline Advisor. AQS0133I will be issued for each connecting Load Balancer instance.

The Network Map on all four BIG-IPs will show the vserver and pool in an Offline state:



Each pool address represents a different site. After the Lifeline Advisor initialization, the workload is placed into a QUIESCED state. Confirm this initial state with a DISPLAY,WORKLOAD,DETAIL report from the primary Lifeline Advisor:

```
AQS0146I WORKLOAD DETAILS
TYPE : ACTIVE/STANDBY
WORKLOAD NAME : WORK1
STATE : QUIESCED
SITE : N/A
SERVERS:
IPADDR..PORT : 16.2.16.200..23
SYSTEM NAME : CHILE SITE : PLEX2 STATE : AVAIL
IPADDR..PORT : 16.11.16.200..23
SYSTEM NAME : RUSSIA SITE : PLEX1 STATE : AVAIL
1 OF 1 RECORDS DISPLAYED
```

In this QUIESCED state, all connection requests to the vserver will be reset. Recall that the Lifeline Advisor controls which site is considered 'Active' for a workload. Activate the workload to PLEX1 by issuing this command on the primary Lifeline Advisor system:

```
MODIFY AQSADV,ACTIVATE,WORKLOAD=WORK1,SITE=PLEX1
```

Lifeline Advisor responds with AQS0153I WORKLOAD WORK1 ACTIVATED ON SITE PLEX1

At this point, the Network Map on all 4 BIG-IPs shows the PLEX1 DVIPA (16.11.16.200:23) to be active:



The primary Lifeline Advisor shows similar information on the DISPLAY,WORKLOAD,DETAIL and DISPLAY,LB,DETAIL reports:

```
AQS0146I WORKLOAD DETAILS
TYPE : ACTIVE/STANDBY
WORKLOAD NAME : WORK1
STATE : ACTIVE
SITE : PLEX1
SERVERS:
IPADDR..PORT : 16.2.16.200..23
SYSTEM NAME : CHILE SITE : PLEX2 STATE : AVAIL
IPADDR..PORT : 16.11.16.200..23
SYSTEM NAME : RUSSIA SITE : PLEX1 STATE : AVAIL
```

1 OF 1 RECORDS DISPLAYED

```
AQS0113I LOAD BALANCER DETAILS
LB INDEX      : 21          UUID      : 66352D6F6D62652D68756C63
IPADDR..PORT  : 66.66.10.53..55173
HEALTH        : 7F          FLAGS     : NOCHANGE PUSH
GROUP NAME    : /COMMON/WORK1-POOL
WORKLOAD      : WORK1
GROUP FLAGS   : CROSS_SYSPLEX
IPADDR..PORT  : 16.2.16.200..23
SYSPLEX       : PLEX2
SYSTEM NAME   : N/A         PROTOCOL  : TCP   AVAIL      : NO
WLM WEIGHT    : N/A         CS WEIGHT : N/A   NET WEIGHT : 00000
FLAGS         : CMQ_DISTDVIPA
IPADDR..PORT  : 16.11.16.200..23
SYSPLEX       : PLEX1
SYSTEM NAME   : N/A         PROTOCOL  : TCP   AVAIL      : YES
WLM WEIGHT    : N/A         CS WEIGHT : N/A   NET WEIGHT : 00001
FLAGS         : DISTDVIPA
... (All 4 BIG-IPs will be displayed)
4 OF 4 RECORDS DISPLAYED
```

Client configuration

Client configuration depends on customer needs. Some applications can allow for a primary and secondary IP address. Both vserver addresses would be configured.

To access via host name, connect to telnet.svt390.com. An F5 DNS solution is highly recommended for this case because it will only return a vserver address to a client requesting the F5 DNS-managed hostname if the vserver is healthy and available.

If the client configuration must be configured with a single IP address, a two vserver solution won't work!

Workload Access via IP Address

Some legacy client applications are only able to access a service via a single IP address. This precludes us from defining a different vserver on each site, because only one vserver can be configured to the client application. When a site fails, client applications would need to be reconfigured to point to the alternate site vserver. To address this problem you will use Route Health Injection (RHI), part of the BIG-IP Advanced Routing Module. You will define the same vserver address in **both** device groups and use routing metrics to pull all client connections into the 'preferred' device group. When a primary device group fails, the secondary device group will now be the lowest-cost route for the vserver and all connections will be handled by the secondary device group.

LTM configuration

Follow the model used for the telnet workload definition but use the same vserver address in each device cluster. Since this vserver address will at times, be routable on either site, you'll use an address from a subnet that is different from either device group's self IP subnet.

First, define the pool for WORK2:

Local Traffic » Pools : Pool List » LL-WORK101-POOL

Properties
 Members
 Statistics

Load Balancing

Load Balancing Method: Dynamic Ratio (member)

Priority Group Activation: Disabled

Update

Current Members

<input checked="" type="checkbox"/>	Status	Member	Address	Service Port	FQDN	Ephemeral	Ratio
<input type="checkbox"/>	◆	16.2.16.200:80	16.2.16.200	80		No	1
<input type="checkbox"/>	◆	16.11.16.200:80	16.11.16.200	80		No	1

Use the same SASP monitor that was created for WORK1.

Create the vserver:

Local Traffic » Virtual Servers : Virtual Server List » LL-WORK101-VS

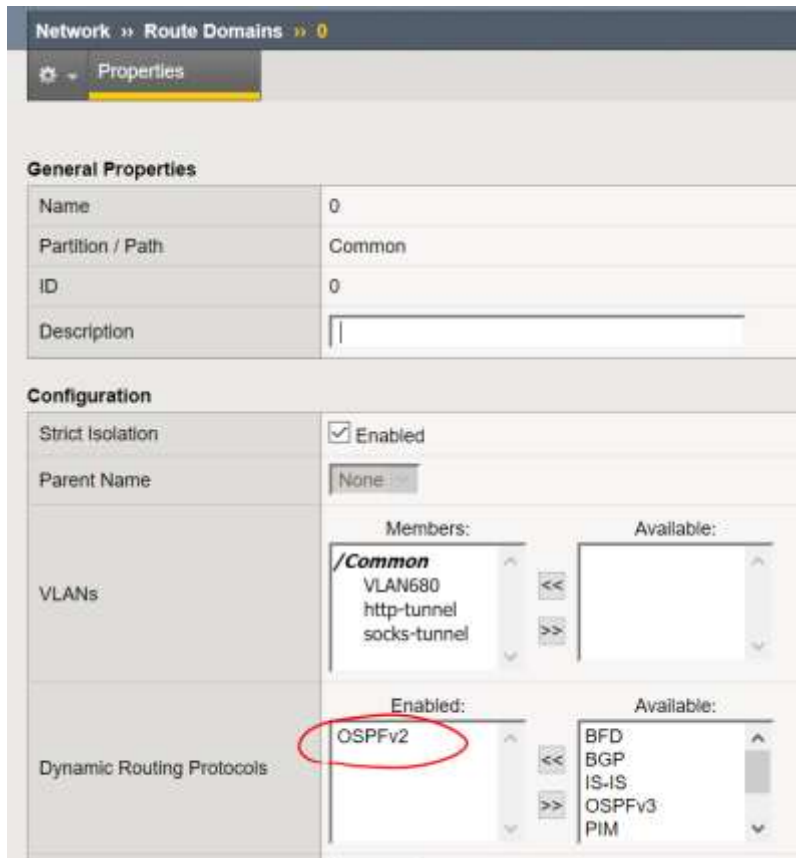
Properties Resources Statistics

General Properties

Name	LL-WORK101-VS
Partition / Path	Common
Description	
Type	Standard
Source Address	0.0.0.0/0
Destination Address	20.20.20.20
Service Port	80 HTTP
Notify Status to Virtual Address	<input checked="" type="checkbox"/>
Link	None
Availability	Offline (Enabled) - The children pool member(s) are down

For Route Health Injection to work, the BIG-IP devices need to participate in a dynamic routing environment. You'll use OSPF for this use case.

On each device, enable OSPFv2 for Dynamic Routing Protocols in the Route Domain:



Create an SSH session to each BIG-IP. Make sure ospfd is running by issuing zebos check:

```
[root@svt156:Active:In Sync (Trust Domain Only)] config # zebos check
=== route domain: 0 ===
nem      is running [13194]
imi      is running [13193]
ospfd    is running [13195]
```

OSPF can only be configured via command line. The **redistribute kernel metric x** parameter allows virtual addresses to be injected into the ospf area. You will specify metric 1 for the BIG-IPs in the device cluster and metric 2 for the secondary device cluster. This panel shows OSPF configuration for a BIG-IP in the secondary device cluster (66.66.11.56):


```

[root@svtf56:Active:In Sync (Trust Domain Only)] config # imish
svtf56.pok.stglabs.ibm.com[0]>enable
svtf56.pok.stglabs.ibm.com[0]#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
svtf56.pok.stglabs.ibm.com[0] (config)#router ospf 1
svtf56.pok.stglabs.ibm.com[0] (config-router)#router-id 66.66.11.56
svtf56.pok.stglabs.ibm.com[0] (config-router)#network 66.66.11.0 0.0.0.255 area 0.0.0.0
svtf56.pok.stglabs.ibm.com[0] (config-router)#redistribute kernel metric 2
svtf56.pok.stglabs.ibm.com[0] (config-router)#exit
svtf56.pok.stglabs.ibm.com[0] (config)#sh run
!
no service password-encryption
!
interface lo
!
interface tmm
!
interface VLAN680
!
router ospf 1
  ospf router-id 66.66.11.56
  redistribute kernel metric 2
  network 66.66.11.0 0.0.0.255 area 0.0.0.0
!
line con 0
  login
line vty 0 39
  login
!
end

```

Next, allow the virtual address to be advertised by OSPF. Take note of the Advertise Route setting – “When any virtual server is available.” This means that when the workload is quiesced on both sites, a host route will no longer advertised in the network. Attempts to ping the virtual address will fail.

Local Traffic » Virtual Servers : Virtual Address List » 20.20.20.20

Properties Statistics

General Properties

Name	20.20.20.20
Partition / Path	Common
Address	20.20.20.20
Traffic Group	<input type="checkbox"/> Inherit traffic group from current partition / path traffic-group-1 (floating)
Availability	◆
State	Enabled
Auto Delete	<input checked="" type="checkbox"/>

Configuration

Advertise Route	When any virtual server is available
Connection Limit	0
ARP	<input checked="" type="checkbox"/> Enabled
ICMP Echo	Enabled
Spanning	<input type="checkbox"/>
Route Advertisement	<input checked="" type="checkbox"/>

Update Delete

Activate workload WORK2 by issuing this command on the primary Lifeline Advisor console:
 MODIFY AQSADV,ACTIVATE,WORKLOAD=WORK2,SITE=PLEX1



Now that a vserver that uses virtual address 20.20.20.20 is available, both clusters inject the host route the network. The primary device group injects the route with a cost of 1 and the secondary device group injects the route with a cost of 2. OSPF processing will prefer the lower cost route. This means that all client connections will be serviced by the active BIG-IP in the primary device group.

From a router in the network, you can see that the metric 1 route is preferred

Routing entry for 20.20.20.20/32

Known via "ospf 226", distance 110, metric 1, type extern 2, forward metric 1
 Last update from 66.66.10.54 on TenGigabitEthernet0/1/0.679, 00:06:57 ago
 Routing Descriptor Blocks:
 * 66.66.10.54, from 66.66.10.54, 00:06:57 ago, via TenGigabitEthernet0/1/0.679
 Route metric is 1, traffic share count is 1

From the client system, open a web browser and load <http://20.20.20.20/> several times. The active device in the primary device group (66.66.10.54 in this case) shows the connections in the vserver statistics:

Statistics » Module Statistics : Local Traffic » Virtual Servers : LL-WORK101-VS

Summary

Display Options

Data Format: Normalized

Auto Refresh: Disabled Refresh

<< Back Clear All Statistics

Traffic Details	Bits		Packets		Connections		
	In	Out	In	Out	Current	Maximum	Total
All	90.1K	108.5K	63	72	7	7	7
Ephemeral	0	0	0	0	0	0	0
Other	90.1K	108.5K	63	72	7	7	7

Force offline both BIG-IPs in the primary device group. The secondary device group route will be preferred. Client connections to <http://20.20.20.20/> will be route through the active BIG-IP in the secondary device group:

This is on 66.66.11.56:

Statistics >> Module Statistics : Local Traffic >> Virtual Servers : LL-WORK101-VS

Summary

Display Options

Data Format: Normalized

Auto Refresh: Disabled Refresh

<< Back Clear All Statistics

Traffic Details	Bits		Packets		Connections		
	In	Out	In	Out	Current	Maximum	Total
All	984	960	3	3	6	6	6
Ephemeral	0	0	0	0	0	0	0
Other	984	960	3	3	6	6	6

Summary

In this example, you saw how BIG-IP redundancy can be implemented across two sites. When a client accesses a workload service by hostname, DNS is used to balance connections across two vservers. When a client must use a hostname, the same vserver is configured on both sites and Route Health Injection is used to direct all connections through a primary BIG-IP device group. The secondary device group is only used in cases where the primary device group is unavailable.