IBM and F5 workload optimization on the new IBM PureFlex System

Creating a highly available, secure, and workload-optimized application delivery environment for the IBM PureFlex System

F5 Networks

IBM Systems and Technology Group ISV Enablement

April 2012

Table of contents

Abstract

High availability and resource optimization are paramount to application delivery and adoption. IT managers and their infrastructures can take advantage of the unique capabilities of the F5 BIG-IP Local Traffic Manager (LTM) and IBM PureFlex System architecture to design a highly available and optimized platform for their mission-critical applications. This paper explains the joint work from IBM and F5 Networks to produce this platform.

Introduction

Applications running across networks encounter a wide range of performance, security, and availability challenges. These problems can cost organizations an enormous amount in lost productivity, missed opportunities, and damage to reputation. IBM® and F5 Networks have worked together to bring to their customers the best solution to solve these problems. By combining the new IBM PureFlex™ System with F5 BIG-IP products, both IBM and F5 are transforming the IT infrastructure into a more agile and flexible infrastructure, both from a network and server perspective.

The BIG-IP product family from F5 is a system of integrated application delivery services that work together on the same best-in-class hardware or virtualized platform. Ranging from load balancing, Secure Sockets Layer (SSL) offload, and web acceleration to network security, application security, access control, and much more, the BIG-IP system creates an agile infrastructure to ensure that applications running on the PureFlex System are always fast, secure, and available.

In this scenario, both BIG-IP hardware and virtualized platforms enhance the service quality of the PureFlex System and provide a high availability environment for an enterprise-class application—in this case IBM WebSphere® Application Server. Using both the hardware and virtual editions of BIG-IP products provide certain advantages in large environments:

- Hardware-based platforms can offload the PureFlex System node's high-processor utilization tasks, such as encryption, compression, and increasing virtual machine (VM) density. It can also be used as a frontline defense from distributed denial of service and other attacks before they reach the application environment.
- The usage of the virtualized edition for individual applications, customers, or lines of business can provide greater control, granularity, and elasticity. Additional modules can then augment standard high availability, optimization, and security scenarios with single-sign on capability, application firewalling, web acceleration, and wide area network (WAN) optimization.
- Hardware and software work in tandem to create a single, unified platform for application delivery, high availability, optimization, and control for PureFlex System environments.

The IBM PureFlex System architecture offers the next generation of IBM technology that brings together server, storage[1], network, and management components in a single chassis. By using a fully integrated system that is managed through a single interface, the system provides faster time to value by reducing system configuration time and simplifying system administration tasks. This system offers a platform for rapid deployment of applications and is designed to handle optimized workloads for both IBM POWER® and Intel® x86 platforms (flexibility). By combining the PureFlex System and the BIG-IP family of products,

[1] Available in future release of the IBM PureFlex System.

customers can enhance the quality of service delivered by the infrastructure, ensuring application availability, security, and optimization for users.

The test environment and plan

The focus of the test is to demonstrate the ability of the F5 BIG-IP system to enhance the quality of service through high availability and front-door services[2] offload of an enterprise application (in this case, WebSphere Application Server) running on the new PureFlex System. To test this, the following environment (as shown in Figure 1) is created.

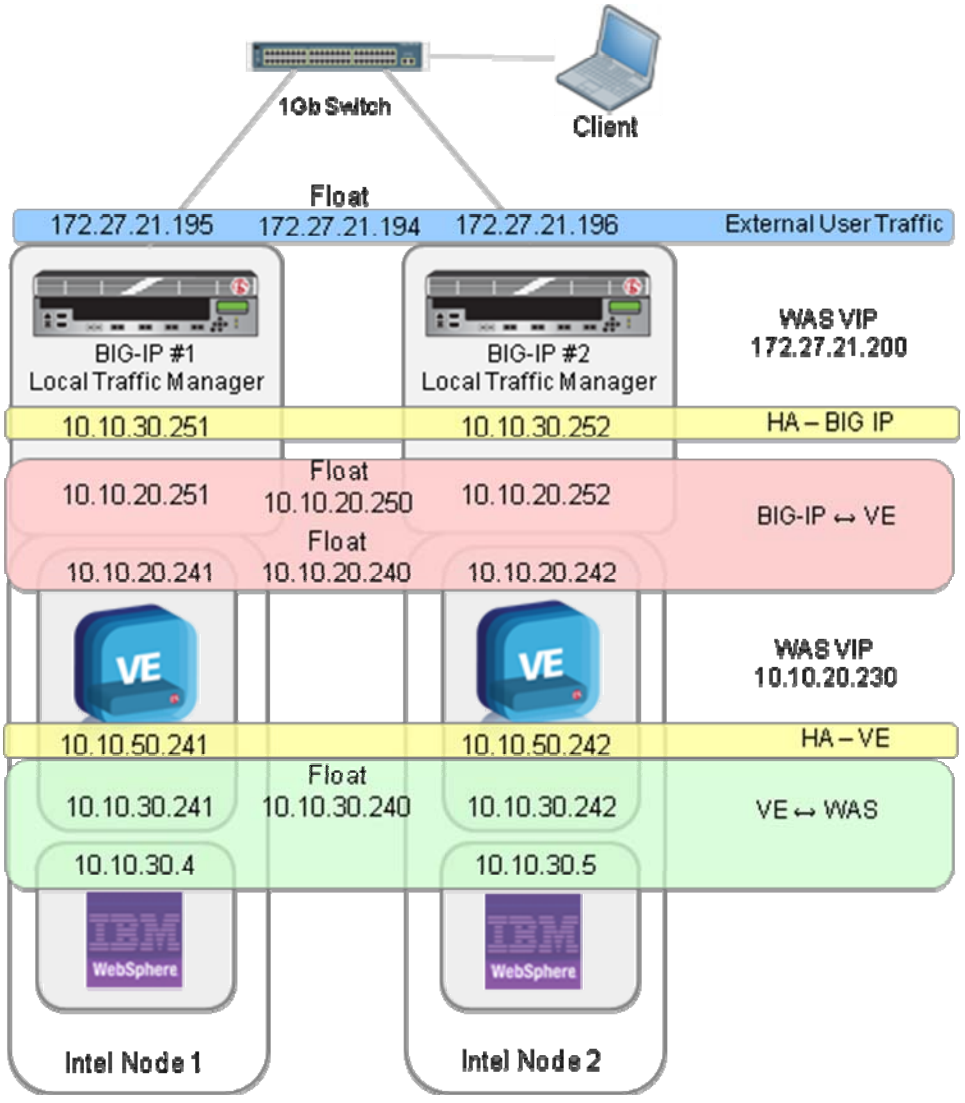

Figure 1: Schematic representation of the testing environment.

[2] Secure Sockets Layer and compression offloading were tested but may also include other services such as prevention of denial of service attacks.

The test environment consists of the following components:

- Two BIG-IP Local Traffic Manager (LTM) 3900 devices, software version 11.1 Build 1943.0
- Two BIG-IP LTM Virtual Edition (VE) instances, software version 11.1 Build 1943.0
- PureFlex System
  - Two Intel nodes
    - 1S 8-core 2.7 GHz Intel Xeon® Processor E5-2680
    - 32 GB RAM
    - Two 300 GB disk
    - One 10Gb Ethernet Emulex One-connect
  - One 10GbE BNT Virtual Fabric Scalable switch
  - VMware ESXi
- IBM WebSphere Application Server V7

The BIG-IP devices were configured in a hybrid scenario, with the physical appliances placed at the edge of the network performing front-door services and two virtual instances performing high availability and application delivery services on the PureFlex System itself. This allows the customer to take advantage of the purpose-built hardware of F5, while still allowing individual application owners or lines of businesses to have granular control of their environment. Each pair of the BIG-IP devices was set up in an active-standby scenario to provide full redundancy of both the physical and virtual pairs.

The PureFlex System was configured with two Intel nodes to provide redundancy of both the WebSphere Application Server environment and the BIG-IP platform.

Test procedures and results

To test high availability of the WebSphere Application server, a web browser client was used to test the status of the application as various failures occurred in the physical and virtualized infrastructure. The team chose the following four different test scenarios:

- Failure of the active physical BIG-IP LTM 3900
- Failure of the BIG-IP LTM VE instance
- Failure of the WebSphere instance running on PureFlex System
- Failure of the entire Intel series node running on PureFlex System

By showing availability of the WebSphere application during all these scenarios, high availability would be demonstrated on the platform.

Test 1: Failure of the active physical BIG-IP LTM device

During this test, the administrator logs in to BIG-IP LTM and performs a shutdown command, while the client actively uses the WebSphere application.

Result: Passed

When the active BIG-IP LTM device was taken down, the standby BIG-IP LTM device came online and serviced the client's request.

Test 2: Failure of the BIG-IP LTM VE instance

During this test, the administrator logs in to the BIG-IP LTM VE instance and performs a shutdown command, while the client actively uses the WebSphere application.

Results: Passed

When the active BIG-IP LTM VE instance was taken down, the standby came online and serviced the client's request.

Test 3: Failure of the WebSphere Application Server

During this test, the administrator logs in to one of the two WebSphere Application Server instances and performs a shutdown command, while the client actively uses the WebSphere application.

Results: Passed

When the WebSphere Application Server instance was taken down, the monitoring capability of the BIG-IP LTM VE instance immediately recognized the failure and started sending traffic to the WebSphere Application Server instance that was up and available.

Test 4: Failure of the Intel series node in the PureFlex System

During this test, the administrator logs in to the PureFlex System platform and performs a shutdown of one of the Intel nodes while the client actively uses the WebSphere Application Server application.

Results: Passed

When the Intel node was taken down, both, the WebSphere Application Server and the BIG-IP LTM VE instance that were running on that node were also taken down. Immediately, the standby BIG-IP LTM VE instance came online and started servicing the requests. The application health-monitoring system of the BIG-IP LTM VE instance (as in test 3), started sending traffic to the instance of WebSphere that was up and available.

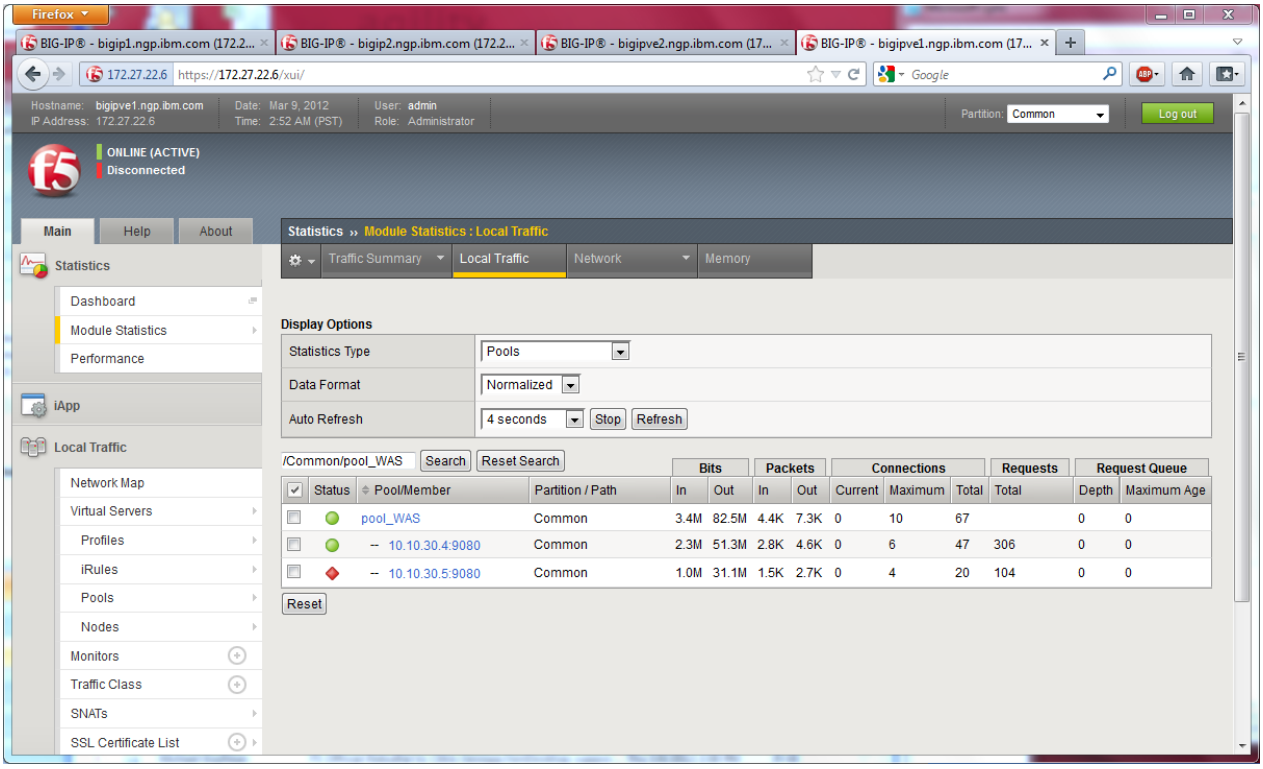

Figure 2: One of the nodes is down while the other takes over as the active node

Best practices

While high availability is a core construct of the BIG-IP system's functionality, the BIG-IP architecture also can reduce server load, optimize, and secure the application. Although not specifically called out in this test environment, the pair of physical BIG-IP devices was set up to terminate the secure sockets layer (SSL) connections from the client and thus offload the encryption process from the WebSphere Application Server. This is one of the best practices when using BIG-IP LTM to provide application delivery services. Figure 3 shows a typical configuration incorporating F5 BIP-IP LTM and the PureFlex System.

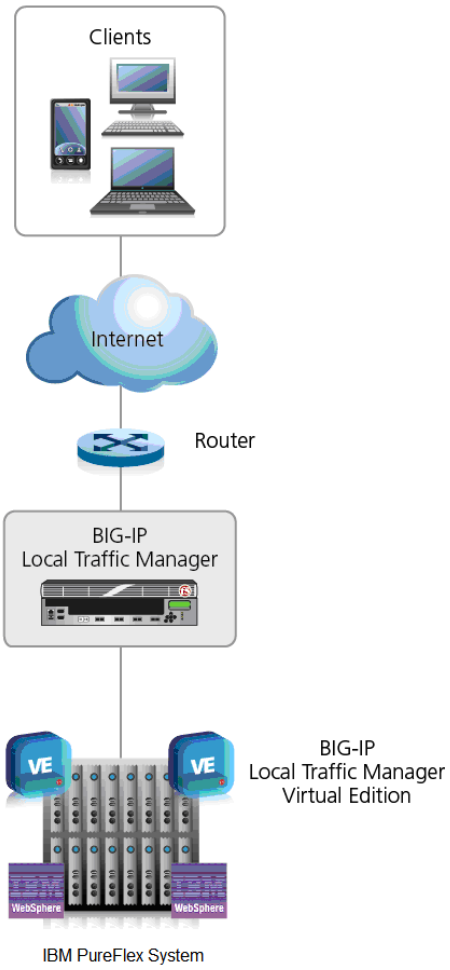

Figure 3 - A typical data center environment incorporating F5 BIG-IP LTM and the IBM PureFlex System

This section lists some of the best practices for your reference:

- Use both static and dynamic load-balancing methods to ensure that the best resources are selected for performance and scale.

- Use specific health monitors to check device, application and content availability, and customized monitors.
- Make use of the device service clustering and configuration synching that goes beyond simple 1:1 active / standby configurations.
- Offload SSL, compression, and other content transformation for effective application integration.
- Use the ability of the BIG-IP LTM to aggregate server-side connections so that they are handled more effectively and increase back-end server capacity.
- Use intelligent caching to offload repetitive traffic from web and application servers.
- If needed, F5 has the unique ability to read the data streams of all IP applications and can switch and persist on information that is unique to a specific vendor's application.
- Turn on intelligent compression to accelerate the application performance up to three times and reduce bandwidth utilization up to 80%.
- If site-to-site data transfer is a requirement, make use of F5 iSession tunnels, the encrypted and accelerated traffic capability of BIG-IP LTM.
- Take advantage of BIG-IP LTM's certification as an ICSA Labs Certified Network Firewall delivering stateful packet inspection to protect data center resources.
- Make use of advanced distributed denial of service (DDoS) protection from over 30 different attack types.
- Turn on resource cloaking.
- Make use of the simple configuration-enabling F5 iApps Templates that provide optimum application-specific configurations in minutes.
- Use iApps Analytics to capture application-specific statistics for visibility, capacity-planning, and troubleshooting.

All of the features listed are part of the core BIG-IP LTM product and are included with standard licensing.

Summary

Ensuring high availability of mission-critical applications is critical for today's enterprise. The F5 and IBM solution helps maintain application availability by providing real-time monitoring of individual PureFlex System instances and the applications that run on them. In addition, BIG-IP LTM can offload and optimize many services from the PureFlex System, helping enhance the quality of service for all the users, increasing overall adoption of the platform and applications.

Resources

The following websites provide useful references to supplement the information contained in this paper:

- IBM Systems on PartnerWorld
  ibm.com/partnerworld/systems

- IBM Redbooks
  ibm.com/redbooks

- F5 BIG-IP Product Family
  www.f5.com/products/big-ip

- F5 Solutions for IBM Applications
  www.f5.com/ibm

- F5 BIG-IP LTM VE Trial Edition
  www.f5.com/trial

Trademarks and special notices

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is

presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.