



Creating a non-clustered HA environment using Big-IP and 3DNS

Amit Budhwani

Clients regularly request for a design to implement an information portal site which does not have transaction or session management but still needs to have high availability without having to invest much in the infrastructure. Designing this kind of solution has always been a challenge for the infrastructure design team. This document will explain a possible solution to design such a system using just a load balancer like Big-IP which is commonly used in such environments.

BIG-IP is an Internet device used to implement a variety of load balancing and other network traffic solutions, including intelligent cache content determination and SSL acceleration and hence is widely used for such sites.

Figure 1 shows the most basic kind of BIG-IP configuration. In it, the unit sits between a router and an array of content servers, and balances the load of inbound Internet traffic across those servers.

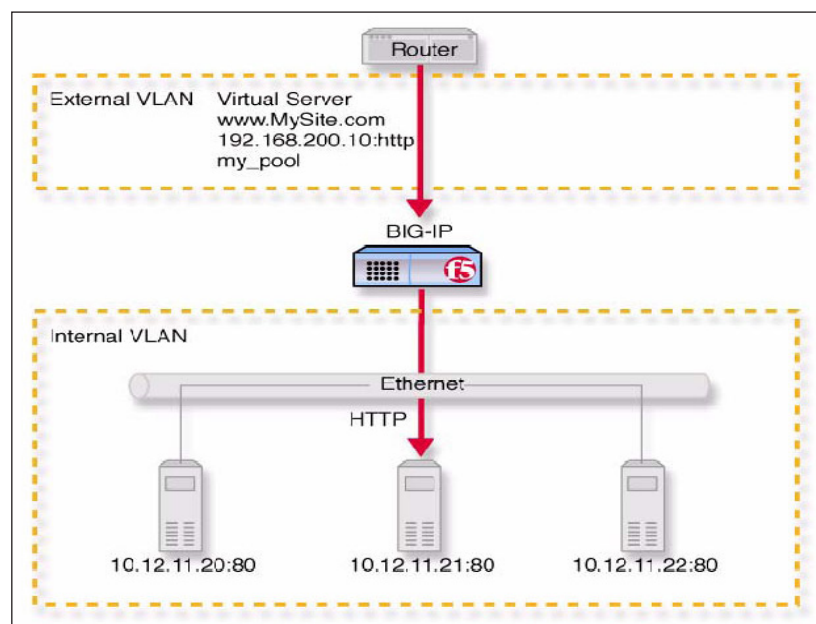


Figure 1: A Basic Big-IP Configuration

The 3-DNS Controller is designed for 24 * 7 availability. It is an intelligent load balancing solution for geographically distributed Internet sites and data centers. It manages and distributes Internet requests across multiple, redundant server sites - regardless of the platform type or combination, and without requiring additional software. End user requests are distributed according to data center and network conditions such as round trip time, packet loss and other QoS metrics, ensuring the highest possible availability for the site.

Through 3-DNS, a policy can be set that allows recovery sites to be used in the event of a disaster. For example, many financial institutions have a primary site in one building, and a disaster recovery site in another. 3-DNS will send all traffic to the primary site, but only if it is performing correctly. If not, traffic is immediately directed to the backup location. 3-DNS eliminates single points of failure, routes traffic away from poorly performing sites and servers, and provides the needed fault tolerance through hardware redundant systems with sub-second recovery times.

Failover is the capability to switch over automatically to a redundant or standby computer server, system, or network upon the failure or abnormal termination of the previously active server, system, or network. Failover happens without human intervention and generally without warning, unlike switchover.

The automation is done using a heartbeat cable that is connected to the two servers. As long as there is a pulse (a.k.a. heartbeat) from the main server to the second server, the second server will not initiate its systems. The second server will immediately take over the work of the first as soon as it detects an alteration in the heartbeat of the first machine. Some failover solutions have the ability to page

or send a message to a pre-assigned technician or center.

Disaster recovery is the combination of processes, policies and procedures of restoring operations critical to the resumption of business, including regaining access to data (records, hardware, software, etc.), communications (incoming, outgoing, toll-free, fax, etc.), workspace, and other business processes after a natural or human-induced disaster.

To increase the opportunity for a successful recovery of valuable records, a well-established and thoroughly tested disaster recovery plan must be developed. This task requires the cooperation of a well-organized committee led by an experienced chairperson.

A disaster recovery plan (DRP) should also include plans for coping with the unexpected or sudden loss of communications and/or key personnel, although these are not covered in this article, the focus of which is data protection. Disaster recovery planning is part of a larger process known as business continuity planning (BCP).

Thus to differentiate difference between failover and disaster recovery, we can say that, failover is more on the machine or network (component) level while disaster recovery is needed when an entire section of the infrastructure goes down in case of natural or unforeseen disaster.

Problem Statement

One of the main criteria of enterprises is to create an environment capable of providing a 24 * 7 system having the least possible downtime. This criterion has always been a challenge for the team involved in setting up the infrastructure and also for the internal technology team. In such cases the solution must address two main goals:

- a. Provide a service with reduced downtime.
- b. Provide a complete failover and disaster recovery scenario which covers the failover for all types of servers.

Big-IP does provide the failover capability out of the box but is restricted to monitoring the heart-beat of the web servers as most installations have a strong firewall between web server and application server preventing communication from Big-IP to the Application Server. The same limitation exists with the database servers and any other backend systems which are integrated in the application.

Currently companies have to spend large sums of money to implement such a solution by using clustering technologies and other HA solutions. These solutions are considered to be over-engineered and expensive when implemented for systems like information portals or low transactional sites.

Solution Description

We will focus on the solution which, without using a highly sophisticated system like clustering, would help us achieve an environment which has the least possible downtime. This solution will help us monitor the heartbeat of the application server, database server and all backend servers using just BIG-IP and additionally provide a failover to a DR site with a 3DNS router and would be ideal for information portals or session-independent sites. This solution will work in any environment having Big-IP and set of any web and app servers without any restriction on technology and OS. Figure 2 illustrates a general infrastructure setup of an organization using Big-IP and 3-DNS.

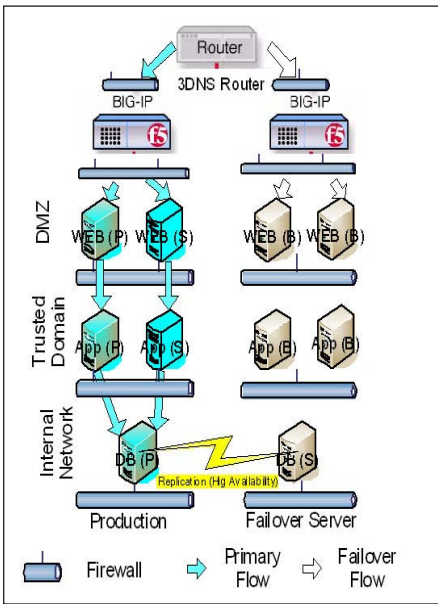


Figure 2: General Infrastructure Design

We have designed a solution which would help Big-IP get access to the application server on trusted domain without having to open a firewall between Big-IP and the application or database server. The steps, as demonstrated in figure 3, are as follows:

- Develop a web page on the application server which fires a query against the database to check if the database connection is successful.
- According to the response from the database, the application returns the status of the server.
- There will be a web page on the web server which will be called by Big-IP, which in turn redirects the request from the above web page to the application server.
- Big-IP, apart from the normal web server heartbeat check, will also be configured to call the web page in the web server as heartbeat check and check the response of the same.
- If the response is unsuccessful, Big-IP will term the entire section to be down and will take this section out of the load balancer list.

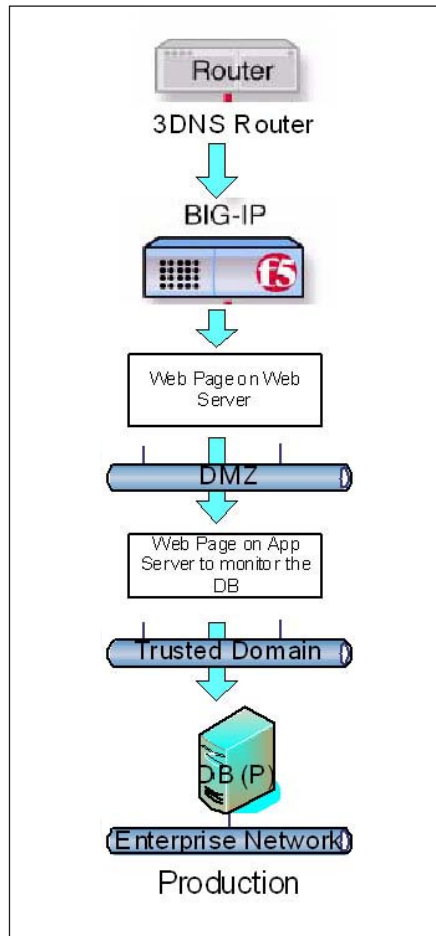


Figure 3: Failover Solution Design

Using the above approach,

- If the database server is down, Big-IP gets the unsuccessful response.
- If application server is down, BIG-IP will get a timeout response from web page call from the web server.
- If the web server is down, BIG-IP will get a timeout from the web server.

In any of the above scenarios, based on above configuration, Big-IP will treat the entire section to be down and will take that section out of the load balancing list. In case of a database or application server failure, the entire environment will be taken out of the loop which in turn will take off the Big-IP from the 3DNS load balancer list resulting in disaster recovery to the DR Environment.

Case Study

One of the world’s largest retail fund houses in Japan commissioned MphasiS to build the next generation customer facing information portal. During the initial study it was determined that 99.9% of uptime (or just 8 hours and 45 minutes downtime per year) would be an extremely important criterion for success. In order to achieve the improved high availability, above solution was deployed which proved successful in meeting the requirements. The setup consisted of 2 web servers (Apache 2.3.3) each having a dedicated connection to one of the 2 application servers (Web-Sphere Portal Server 5.1.2). All these systems were built on the Linux environment running on Red Hat 4 (all Servers except web servers) and Red Hat 5 (only web servers).

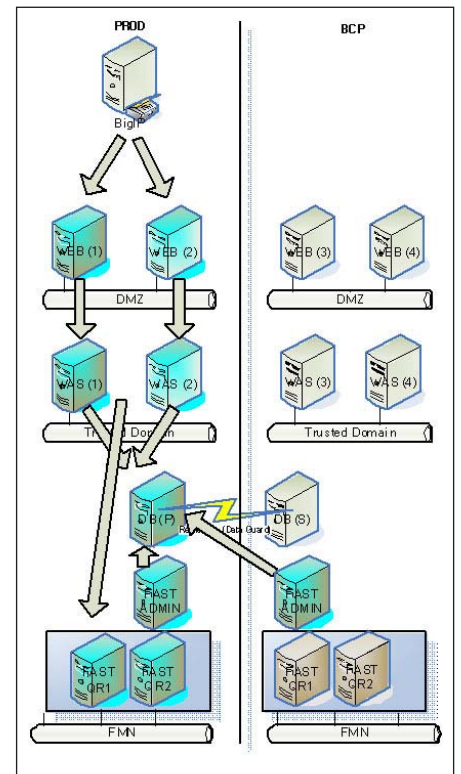


Figure 4: Design as implemented by client

The solution consists of a CGI file on the web server which is called by Big-IP. In turn, the CGI file calls a JSP file on the portal server which checks the connection to the database server and returns a

success message back to the CGI script. This CGI is monitored at an interval of 5 seconds which proves to be a great success in achieving the targeted result without investing in additional clustering technology.

Figure 5 explains the scenario that was achieved by the above solution.

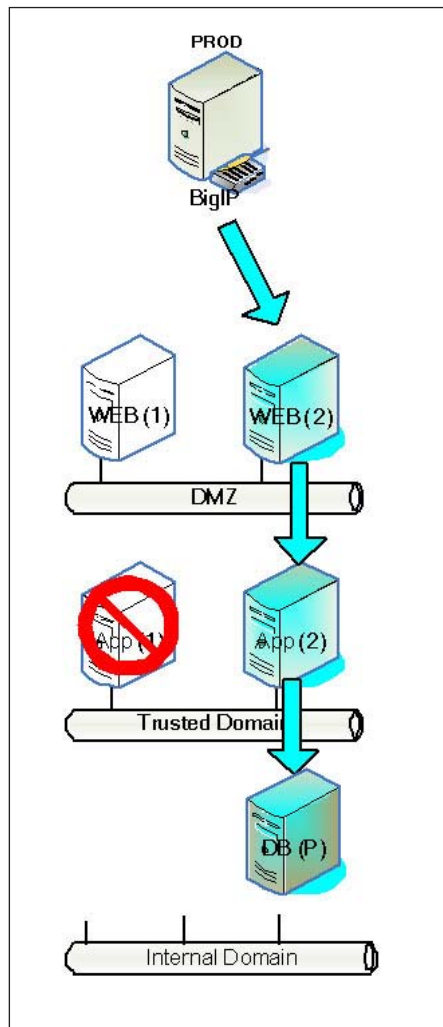


Figure 5: Failover Design in case of Application Server Failure

Conclusions

Companies had to choose between heavy clustering technologies which involved high costs and the ability to implement high availability environments. In such cases many companies chose to forgo the HA environment, thus reducing the overall business value of information sites and low transaction sites which did not depend on session level information. With our solution companies can continue to deploy non-session dependant sites in a high availability environment without having to invest in technologies such as clustering. The solution is non-invasive and hence can be de-ployed on existing infrastructure which already includes Big-IP, thereby providing greater business value.

About Mphasis

Mphasis supports G2000 companies around the world in the improvement of their business processes. Our unique strength lies in the intersection of our Information Technology (IT) and Business Process Outsourcing (BPO) capabilities. While our forte is in the BFSI (Banking, Financial Services, Insurance) and Technology industries, our focus extends to the Healthcare, Mobile, Logistics, Telecom, Life Sciences, Consumer Electronics and Utilities industries. The convergence of technologies such as web services, workflow software and business performance monitoring along with business intelligence and customer focus drive all our services delivery offerings.

For more information on a Pilot Project option, write to vtr@mphasis.com

Contact us

USA

Mphasis
460 Park Avenue South
Suite # 1101, New York
NY 10016, U.S.A.
Tel: +1 212 686 6655
Fax: +1 212 686 2422

UK

Mphasis
100 Borough High Street
London SE1 1LB
Tel: +44 20 30 057 660
Fax: + 44 20 30 311 348

Mphasis
Edinburgh House
43-51 Windsor Road
Slough SL1 2EE, UK
Tel: +44 0 1753 217 700
Fax: +44 0 1753 217 701

INDIA

Mphasis
Bagmane Technology Park
Byrasandra
C.V. Raman Nagar
Bangalore 560 093, India
Tel: +91 80 4004 0404
Fax: +91 80 4004 9999