A Link Load Balancing Solution for Multi-Homed Networks

Overview

An increasing number of enterprises are using the Internet for delivering mission-critical content and applications. By maintaining only one link to the public network, they are faced with a single point of failure and serious network vulnerability. When link failures occur, all of the investments that an enterprise has made in redundant systems are immediately rendered useless, and the organization and its customers are left with significant downtime and financial loss.

Some organizations consider multi-site deployments to address their availability concerns. However, small to mid-size enterprises quickly find this option cost prohibitive. Multi-site deployments not only face the increased technical complexity of hosting distributed Web applications (such as synchronizing backend data), but require costly expenditures for additional equipment, rack space, software and personnel.

Traditionally, enterprises chose to address their network reliability issues by installing a secondary public network connection to a site. This approach, called multi-homing, typically used Border Gateway Protocol (BGP). However, this deployment provides limited route control, increases latency, and requires significant administrative overhead.

The BIG-IP® Link Controller provides a simpler, less costly way to increase reliability for sites with multiple network connections. By using the BIG-IP Link Controller, organizations can build a reliable, fault tolerant, and highly available multi-homed network.

This white paper examines the challenges faced by enterprises when deploying a multi-homed network with BGP, and how the BIG-IP Link Controller provides a solution to increase the availability and performance of Internet connectivity, while greatly simplifying deployment and management.

Challenge

To date, organizations have primarily relied on BGP as a means for directing traffic over multiple Internet links. BGP, an inter-domain routing protocol, was designed to enable IP routers to direct packets traversing along the Internet from point A to point B. While BGP is a core technology for routing, implementing multi-homing using BGP can be extremely difficult and does not provide a proper mechanism for ensure dynamic, flexible routing based on link performance, corporate policy or utilization cost models. Specifically:

BGP is costly and complex to deploy - Multi-homing using BGP is extremely complex, and requires hard-to-find network expertise that organizations often do not have. Some solution scenarios require organizations to purchase costly, high-end routers that are capable of running BGP with sufficient memory to maintain its corresponding routing table, designated IP address blocks, and Autonomous System Number (ASN), which is a globally unique routing domain identifier.

BGP introduces latency and performance issues - With BGP, gateway hosts exchange routing information that are used to create a view of the Internet to base routing decisions on. When a change is received, such as a new or failed route, the router initiates an incremental update to peer routers notifying them of the change. Those routers update their table and then propagate the change to their peer routers using the same process. This process, known as convergence, continues until the information between all the routers has been synchronized. The time required for convergence is lengthy and can take up to 30 minutes to complete. In the interim, traffic delays and misrouted traffic can occur, adversely affecting the ability of customers to utilize network and application resources.

BGP requires ISP cooperation and maintenance - In many parts of the world, organizations are unable to obtain a global ASN. This requires them to lease a private ASN and IP address block from one of their ISPs in order to use BGP. Both ISPs must then agree to provide routing for the address block and advertise the ASN as their own. This solution scenario requires

tremendous cooperation between the organization's competing ISPs, which can limit an organization when initially choosing a secondary ISP. Even after the relationship is established, when problems occur with the configuration, organizations are forced to wait while the ISPs try to determine the cause of the problem, and who is responsible.

BGP provides inferior traffic management capabilities - BGP route selections are based on hop count. This greatly limits the flexibility of the implementation and decreases the accuracy of the routing decision it makes to ensure the best possible link selection. BGPs lack of granular control means it can not identify how traffic should be dynamically routed based on line saturation, performance, or cost.

BGP lacks link flexibility - Because of the size of BGP's routing table and frequency of updates, the ability of an enterprise to use most inexpensive, widely available link connections such as DSL, Cable, or Wireless, is eliminated. BGP's overhead locks organizations into using larger bandwidth links at greater recurring costs.

Solution

The BIG-IP Link Controller provides high availability for enterprise sites with multiple Internet links and ISP providers. It monitors the health and availability of each connection and transparently directs inbound and outbound traffic to the best available WAN gateway to minimize user response times and bandwidth costs. The BIG-IP Link Controller works regardless of connection type or provider and eliminates the need for ISP cooperation, designated IP address blocks, ASN, and other deployment barriers associated with multi-homing using BGP. It extends the ROI for corporate connectivity by ensuring traffic is optimally managed over existing bandwidth investments.

Simplified multi-homing deployment

The BIG-IP Link Controller simplifies multi-homing deployments through its flexible and comprehensive design. Organizations no longer need ISP cooperation, large bandwidth connections, designated IP address blocks, ASNs or high-end routers to protect themselves from network failures. Because the Link Controller solution does not rely on BGP to provide fail-over capabilities, the solution does not suffer from the same problems as multi-homing with BGP does, such as latency, high update overhead, and inferior traffic management. Using the Link Controller, organizations benefit from guaranteed availability without delays or costly misrouting, the ability to use inexpensive bandwidth solutions, and complete granular control for link selection based on performance, costs, and business policies.

Intelligent traffic management

The BIG-IP Link Controller can be easily deployed into any network, and offers organizations a turnkey solution that is transparent to users and ISPs, and provides superior traffic management functionality to BGP. Built on the successes of F5's award-winning traffic management products, the BIG-IP Link Controller provides comprehensive intelligent traffic switching for both inbound and outbound traffic.

Secure Network Address Translation Automap - Secure network address translation (SNAT) automap provides a secure mechanism for translating internal non-routable IP addresses into routable addresses. The BIG-IP Link Controller uses SNAT automap to dynamically translate the IP address of an internal resource into one that is associated with the best link. For instance, if the best link to direct an outbound request is link A (based on the configured load balancing criteria), the BIG-IP Link Controller automatically translates the packet's source IP from the internal resource to the Link Controller's self IP address for Link A. This ensures that the response to the request is returned through the same link and ISP.

Wide IP - The Link Controller load balances inbound connections by using wide IPs. A wide IP is a collection of one or more domain names that maps to one or more groups of virtual servers managed by the Link Controller. The Link Controller load balances

name resolution requests across the virtual servers that are defined in the wide IP that is associated with the requested domain name.

Ensuring Availability

The BIG-IP Link Controller guarantees reliable network connections and eliminates downtime by detecting any type of connection outage. In the event of an ISP or link failure, traffic is transparently and dynamically directed across other available links to provide seamless traffic flow. It allows an organization to provide complete reliability for a site by shifting IP traffic away from failed or over-saturated Internet or leased-line connections -- whether caused by networking errors (such as router failures), or by a disruption in the service an ISP provides.

The BIG-IP Link Controller provides several levels of monitoring to ensure link and ISP outages are detected quickly:

Default Gateway Monitoring - The Link Controller automatically creates ICMP monitors that check availability and detect the performance of critical gateway routers for a given link.

Transparent Monitoring - Transparent monitors enable the Link Controller to check the availability of a device in an ISP network or on the Internet to determine the availability of the entire link. These monitors are used to detect errors that occur beyond the gateway router, such as errors like ISP software failures that might lead to a loss of connectivity.

Rate Monitoring - Rate Monitors examine the performance of links based on the traffic that passes through the BIG-IP Link Controller. For some sites, this eliminates the need for external performance monitoring, increasing the speed and efficiency of tracking performance metrics for a given link.

Increase scalability and link performance

The BIG-IP Link Controller enables flexible scalability while reducing cost for a site. Traffic can be routed based on performance, bandwidth cost, and bandwidth availability using metrics including usage limits and the pricing structure of purchased bandwidth. This enables customers the flexibility to create a single *virtual link*, utilizing the total bandwidth of all their connections. It also provides bandwidth management flexibility across different size links and connection types, preventing bandwidth bottlenecks while minimizing the costs of inefficient bandwidth utilization. The Link Controller allows customers to aggregate lines to provide lower-cost bandwidth redundancy while minimizing spending on dark fiber and idle lines.

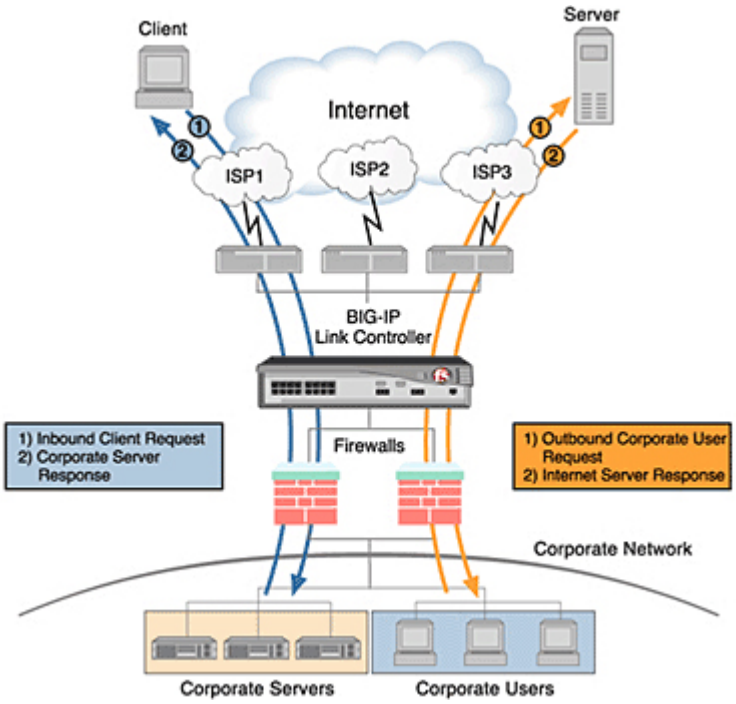

***Figure 1**: The BIG-IP Link Controller provides comprehensive intelligent traffic switching for both inbound and outbound traffic to a data center*

Static routing - Static load balancing modes distribute traffic based on configured metrics and are not influenced by dynamic environmental variables.

Round Robin - Round Robin mode is a static load balancing mode that bases connection distribution on a set order. Round Robin mode sends a connection request to the next available server for inbound traffic and a link for outbound traffic in that order.

Ratio - The Ratio load balancing mode distributes connections across an array of virtual servers in proportion to the ratio weights assigned to each individual virtual server for inbound traffic and across the configured links for outbound traffic.

Link Cost load balancing - This defines the price you pay for bandwidth per megabit per second. By using Link Cost load balancing, the BIG-IP Link Controller routes each request across the least expensive link to minimize bandwidth cost across varying priced lines.

Variable Cost load balancing - This defines multiple price points for a single link in order to perform lowest cost load balancing on links with variable cost thresholds. (For example, a burstable T1 line may cost $800 up to 1 Mbps / month and then jump to $1000 for traffic over 1 Mbps.)

Policy Based Traffic Control - The BIG-IP Link Controller provides flexible traffic control to define segments of traffic and designate them toward specific resources. Source, Destination, Application Switching enables customers to direct portions of traffic over designated links according to business policies. High priority traffic, such as FTP or email, can be routed over a single preferred connection, while lower priority traffic is routed over an alternate link.

Dynamic routing - The BIG-IP Link Load Balancer provides flexible traffic management across multi-homed networks by allowing customers to weight the importance of link bandwidth, performance, and cost to distribute traffic according to their needs. Dynamic load balancing modes base the distribution of requests on live data, such as current server performance and current connection load.

Quality of Service - The Quality of Service load balancing mode is a dynamic inbound load balancing mode that bases connection distribution on a configurable combination of the packet rate, completion rate, round trip time, hops, virtual server capacity, kilobytes per second, and topology information.

Dynamic Ratio - Dynamic Ratio mode for outbound traffic is like Ratio mode, except that ratio weights are based on continuous monitoring of the links and are, therefore, continually changing.

Evaluate ISP and link performance

In choosing an ISP provider, enterprises typically have few means to evaluate the quality of the connectivity provided relative to other providers. Evaluating the real service received from an ISP is important, and would help many organizations in evaluating their ROI for their connectivity.

The Internet Link Evaluator - The Internet Link Evaluator provides a unique view of the network by measuring the total performance that an ISP passes along to those accessing a site. By measuring Round Trip Time, Completion Rate, and average router Hops for each link, businesses receive a holistic view of Link and ISP performance. This information can be used to make critical business decisions that enable cost savings and performance gains as a result of choosing the best service providers.

With the Internet Link Evaluator, organizations can:

- Spot degradation in link performance / troubleshoot performance issues over ISP links
- Evaluate an ISP provider's ability to quickly serve your Internet users (Round Trip Time)
- Evaluate a provider's connection quality from your network to your users (Completion Rate)

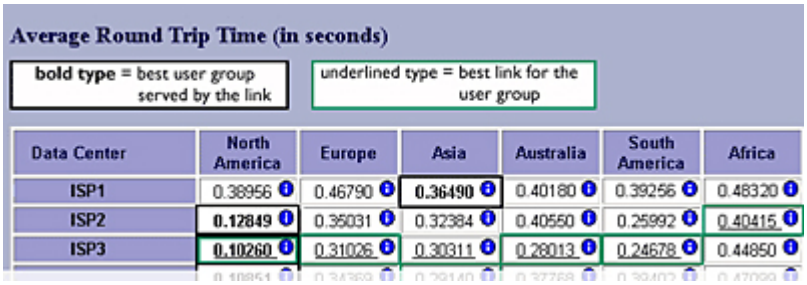

***Figure 2**: The Internet Link Evaluator*

Open/Upgradeable Solution

Enterprises are looking for a solution that will consolidate functions in a cost effective, easy to manage product - one with an extensible platform that can grow with an organization's changing needs. The BIG-IP Link Controller offers organizations an extensible platform:

- Integrated firewall load balancing for high availability to redundant firewall deployments
- Upgradeable to include enhanced local load balancing of servers and caches
- Upgradeable to include comprehensive, multi-site global load balancing and disaster recovery
- Port dense, high performance platform that scales with an organization
- Supports multi-gigabit throughput and any number of ISPs