

Conquering Multi-Homed ISP Link Challenges

Businesses are still struggling to deliver continuous application high availability to their customers. In fact, a report from Infonetix states that service provider outage contributes to 30% of the downtime businesses experience.

Small and medium businesses must also provide high availability for their applications, but often don't have the IT budget to fund a second data center. A workaround for this dilemma is a multi-homed solution for ISP links, side-stepping the hurdles of using the Border Gateway Protocol (BGP).

Businesses that are using a multi-homed network or thinking about it are also faced with the challenges that are inherent to a new class of applications such as OWA, SIP, and Web Portals, etc. These applications are putting immense pressure on ISP links. Businesses have to not only worry about link availability, but also about application performance and how well their ISP links are being utilized. Problems such as bandwidth oversubscription, contention, inefficient link utilization, and inter-ISP routing issues are rampant and costing businesses a lot of revenue and customers while increasing their operational expenses.

This white paper explores the challenges that businesses face when deploying a multi-homed network and how the BIG-IP® Link Controller provides a cost effective solution to not only manage availability, but also to optimize the performance of ISP connectivity in an operationally efficient manner.

How do I multi-home my ISP links?

Today, businesses continue to lose millions of dollars because of ISP availability problems. However, deploying a cost effective solution that is operationally efficient can also be a challenge. Some alternatives to multi-homing networks include:

- **Multi-home with two links and one ISP link** – Businesses can multi-home their sites with two links from the same ISP. This is a cheaper solution to deploy, but not very efficient because an outage at the ISP creates a bottleneck as both links become unavailable.
- **Multi-home with two links and two ISP links** – Typically, businesses multi-home their sites with two links from two separate ISPs and use BGP to route across them. Although BGP can provide link availability in the case of a failure, it is a complex routing protocol that is not well suited to provide multi-homing and intelligent link routing. BGP is costly and complex to deploy because it needs special Autonomous System Numbers (ASN) from the ISPs and router upgrades to be installed. In the case of a failure, ISP cooperation is needed sometimes for link recovery, and BGP propagation is notoriously slow for the different routes to converge. This causes long and unpredictable failover times and does not meet the high availability requirements for a business.

How do I optimize my multi-homed links?

Businesses are realizing that multi-homing ISP links isn't enough. They are continuously plagued with link performance issues and usually add more bandwidth to solve the problem. This creates an endless cycle of investment in bandwidth with diminishing returns. To make things worse, organizations are experiencing an unprecedented demand for bandwidth to deliver IP services, but are unable to cost-effectively and intelligently manage bandwidth.



How do you strike a balance between using bandwidth efficiently and continuing to add bandwidth to solve performance problems? What often happens is a tradeoff between new application deployment and bandwidth usage while trying to avoid oversubscribing. These issues include:

- **Application prioritization** – Organizations are plagued with prioritizing the different types of traffic traversing their saturated WAN link. They are often at the mercy of bandwidth-hogging applications that degrade performance and affect productivity. What they would like to do is to prioritize traffic based on application type, user, location, and link so that high-priority or critical applications always get the bandwidth they need.
- **Bandwidth oversubscription** – Organizations are constantly confronted with excessive bandwidth usage over the subscribed limits. As bandwidth becomes scarce, application performance starts suffering and ultimately affects the profitability of the business. With the increase in the number of client types accessing their web site, businesses can no longer use the one-size-fits-all bandwidth management policy. What they would like to do is to create custom bandwidth policies for each type of user to deliver a superior end-user experience.
- **TCP inefficiencies** – TCP/IP is a protocol that was originally developed for the LAN. When run over the WAN, TCP has a lot of built-in inefficiencies being a connection-oriented protocol, which causes poor application response times. TCP inefficiencies generate a lot of unnecessary chattiness and higher recovery times for lost packets, dramatically affecting the bandwidth utilization of the link. Clients with longer physical distances are also affected as TCP reduces raw throughput and causes the WAN link to be underutilized. Dial-up clients suffer longer latencies and oversubscribe the link while broadband clients are less able to fully fill the WAN link.
- **Link management** – Lack of visibility into link performance and availability, the inability to create policies to manage quality of service, and lagging support for next-generation networks greatly increases management overhead and total cost of ownership. ISP peering problems are also becoming prevalent. Peering disputes among major ISPs are causing serious connectivity problems for businesses, translating into loss of revenue and profitability. The peering disputes between the various tiered providers causes some serious connectivity problems.
- **IPv6 deployment** – IPv6 deployment has already begun, and service providers have started offering IPv6 services. In fact, the Department of Defense has mandated that all their networks will run IPv6 by 2008. Businesses are preparing for forklift upgrades to their existing networks to deploy and access IPv6 services.

The BIG-IP Link Controller

F5's BIG-IP Link Controller is integrated with the BIG-IP Local Traffic Manager and built on the TMOS architecture that provides a holistic way to solve your WAN link challenges, including performance, security, high availability and management. The BIG-IP Link Controller provides you with the flexibility and scalability you need to adapt to evolving bandwidth and ISP high availability needs while delivering the best quality of service to your end users.

Link high availability without the complexity

The BIG-IP Link Controller simplifies multi-homed deployments so you no longer need ISP cooperation, large bandwidth connections, designated IP address blocks, ASNs, or high-end routers to protect your network from ISP failures. Using DNS-based technology that removes the dependency on BGP to provide failover capabilities, the BIG-IP Link Controller eliminates multi-homed problems such as latency, high update overhead, and inferior traffic management. With the BIG-IP Link Controller, you can benefit from guaranteed availability without delays or costly misrouting. It also gives you the ability to aggregate inexpensive links, with more granular control over which link to use based on performance, costs, and business policies.


Taking link performance to a new level

F5’s BIG-IP Link Controller is the only product to deliver integrated, link-specific optimization services to solve link performance problems in an operationally efficient manner. The BIG-IP Link Controller can help boost bandwidth savings, improve application performance, and eliminate the need for additional bandwidth. The BIG-IP Link Controller does this using the following optimization services:

- Rate Shaping to prioritize WAN link bandwidth for critical applications**

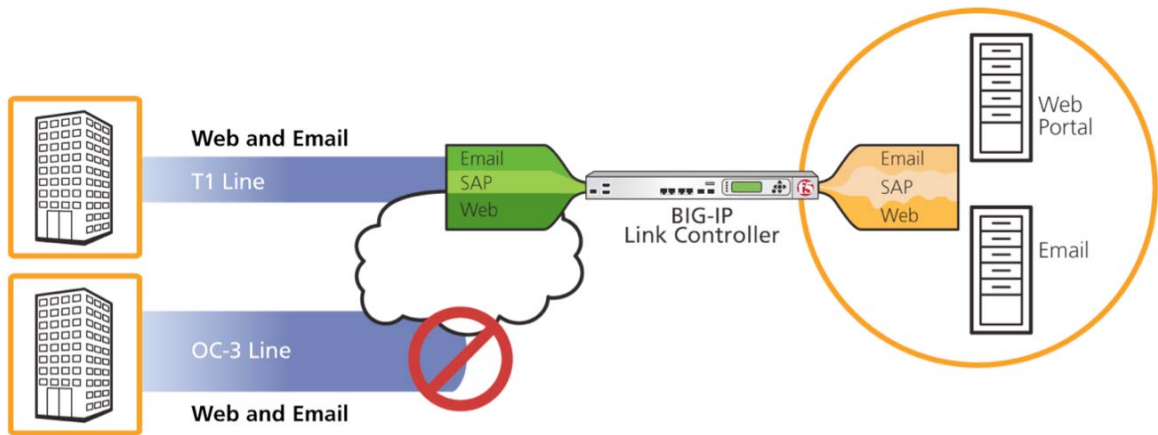
With Rate Shaping, you can define traffic and application limits, control the rate at which those resources are allowed to spike or burst, provide queuing to prioritize traffic types, and define relationships where certain traffic types can borrow from other traffic types. Using the BIG-IP Link Controller, you can control available bandwidth so that higher priority and latency sensitive applications such as VoIP get the bandwidth they need to perform well. You can also manage bandwidth usage according to business policies such as application type and type of user to avoid contention for bandwidth.

In the following example, an online business has two ISP links that use BGP to provide high availability to their site. Their main applications are Web Portal and email. When the primary link is available, everything runs smoothly and applications perform well. However, when their primary ISP link went down and the backup ISP link became active, their web application consumed all the available bandwidth, starving their email application. Since 60% of their business was done via email, there was a significant impact on revenue.



Web Application Hogs Bandwidth, Starves Email

Using Rate Shaping, the BIG-IP Link Controller can re-prioritize the web and email applications when the back up link is active. By carving out bandwidth for each application, the BIG-IP Link Controller ensures that the applications don’t contend for the same bandwidth and get the bandwidth they need to perform well.



Web Application and Email Get the Bandwidth They Need

- Compression to reduce WAN Link bandwidth consumption**
 An optional compression module gives you the ability to compress web application traffic, reducing bandwidth usage without buying more bandwidth. F5 provides a very flexible and tunable compression engine so you can configure compression based on Document type, Traffic type, and Network conditions such as Round trip time. Compressing HTTP traffic helps you save on monthly ISP bills, reduce bandwidth usage, and accelerate end-user performance.
- TCP optimizations over the WAN**
 TCP protocols are extremely chatty in nature and can easily consume unnecessary bandwidth. At the heart of the BIG-IP Link Controller is TMOS with a highly optimized TCP stack to reduce TCP inefficiencies. Using sophisticated TCP WAN optimization techniques such as Fast Retransmits and Recovery, Adaptive Congestion Windows, Selective Acknowledgements, and Congestion Notification, F5's TCP stack reduces latency due to retransmissions of lost packets, distance and network congestion, and improves end-to-end performance by providing a significant increase in bandwidth. Using specific utilization improvements such as TCP Slow Start with Congestion Avoidance, Bandwidth Delay Control, and Timestamps, organizations can realize a significant gain from efficient bandwidth utilization on their existing links for clients over long distances.

Managing multiple links with ease

The BIG-IP Link Controller provides a powerful yet simple way to manage multi-homed ISP links with link object management, real time and historical views of link bandwidth usage, and multiple health monitors to accurately determine the availability of a link. The BIG-IP Link Controller provides other management tools including:

- Topology-based routing to the best ISP link**
 With the BIG-IP Link Controller, you can route traffic over multiple links based on the location of the user. It does this by using its topology database and IANA tables to determine the location of the user, and then routes the traffic over the desired link based on pre-defined policies. This way, you can choose the best performing link for the best end-user experience based on the location of the user. This also helps solve ISP Peering disputes that result in a loss of connectivity by providing a multi-homed alternative and ensuring that users get directed to the right ISP link.
- iRules for programmatic control**
 You can route traffic over multiple links based on TCP/IP parameters such as Source IP, Destination IP address, port, etc. With iRules, a TCL-based scripting language, you can quickly define policies such as distribute traffic over links based on application type,



Quality of Service, and client types, to improve application performance. The following iRule dynamically measures the Round Trip Time for each link and selects the link with the lowest value (best QoS) for SIP applications.

```
rule rule1 {
  when CLIENT_ACCEPTED {
    Log local0. "RTT for 10.0.0.150 = [ROUTE::rtt "10.0.0.150"]"
    log local0. "RTT for 12.0.0.150 = [ROUTE::rtt "12.0.0.150"]"
    if [TCP::local_port] == 5060 { # SIP Traffic
      if {[ROUTE::rtt "10.0.0.150"] < [ROUTE::rtt "12.0.0.150"]} {
        use pool gateway_pool member 12.0.0.150:0
      } else {
        use pool gateway_pool member 10.0.0.150:0
      }
    } else {
      use pool gateway_pool
    }
  }
}
```

- **Support for next generation IPv6 networks**

Using the BIG-IP Link Controller's IPv6 gateway, you can host IPv6 services while still providing access to IPv4 customers, and translate between them without burdening the network.

Conclusion

The BIG-IP Link Controller combines the best-of-breed link optimization and availability services to give you maximize availability, performance, and manageability for multi-homed ISP links. Integrated with the BIG-IP Local Traffic Manager and built on TMOS, the BIG-IP Link Controller evolves with your changing bandwidth needs to deliver:

- Reliable WAN connections for business-grade Internet connectivity
- Efficient WAN link bandwidth with compression and rate shaping
- Improved link performance with an optimized TCP/IP stack
- The highest quality of service and speed by directing traffic directed over the best link and ISP
- Better ROI by aggregating inexpensive links
- Lower multi-homed costs via BGP

About F5 Networks

F5 Networks is the global leader in Application Delivery Networking. F5 provides solutions that make applications secure, fast, and available for everyone, helping organizations get the most out of their investment. By adding intelligence and manageability into the network to offload applications, F5 optimizes applications and allows them to work faster and consume fewer resources. F5's extensible architecture intelligently integrates application optimization, protection for the application and the network, and delivers application reliability – all on one universal platform. Over 16,000 organizations and service providers worldwide trust F5 to keep their applications running. The company is headquartered in Seattle, Washington with offices worldwide. For more information, go to www.f5.com.