

MULTI-HOMING
WITHOUT BGP

“...Enterprises lose \$700 billion from downtime; losses fall into three categories: lost revenue (17%), lost productivity (73%), and cost associated with fixing problems (5%)...”.

IHS

Overview

INTERNET CONNECTIONS are becoming faster and more economical, yet at the same time Internet uptime and network latency continue to be a problem. If you have ever had an outage at your ISP that has affected your business, you understand.

Because the Internet is becoming cheaper AND more mission critical, businesses that could only afford a single T1 or DSL years ago are now open to aggregating two or more broadband connections to the Internet from different ISPs (multi-homing) if they can find a simple, inexpensive way to make it happen.

ISPs and large enterprises have addressed the multi-homing problem for years by using Border Gateway Protocol (BGP) to connect multiple Internet backbones. Unfortunately, the cost and complexity of this BGP routing does not lend itself to small to medium businesses. Also, many broadband, DSL and fixed wireless providers will not support BGP to their end users.

Ecessa overcomes these restrictions by using Network Address Translation (NAT) and Dynamic Domain Names Service (DNS) to direct each new TCP session. Ecessa solutions not only load balance users on the corporate LAN trying to get out, but they also load balance users from the Internet trying to access web and email servers hosted on the corporate LAN. All of this happens while still simultaneously providing automatic ISP/link fail-over for an always available Internet connection.

Major Barriers with BGP

Previously, small and medium sized businesses that needed to ensure reliable Internet connectivity simply didn't have a viable solution. These organizations typically attempted to address this problem through BGP (Border Gateway Protocol), but would run into significant deployment barriers that virtually eliminated it as a viable solution. Here's why:

- BGP is costly and complex to deploy
- BGP introduces latency and performance issues
- BGP requires ISP cooperation and maintenance
- BGP requires difficult to obtain collections of Internet address blocks (ASNs)
- BGP provides inferior traffic management capabilities

Differences Between ECESSA and BGP

Not only is multi-homing with BGP extremely difficult to setup and maintain, it does not optimize WAN link performance. This is because the burden is placed upon the ISP to influence the BGP tables. BGP based solutions increase the number of routing table updates on all routers between the links in a multi-homed network. This negatively impacts the performance of the routers and the ability to truly offer optimized traffic management.

Further, all BGP based solutions encourage announcing the same CIDR block (Classless Interior- Domain Routing) to multiple ISPs. This directly increases the number of BGP global routing table entries, further impacting performance, and discouraging ISPs to cooperate (because in general only one ISP can aggregate a given CIDR block). In contrast, the Ecessa approach does not require announcing the same CIDR block to multiple ISPs, thereby having no effect on the global routing table size or routing performance.

BGP – An Obsolete Approach

As mentioned earlier, there are different types of failover, some that are not entirely automatic by intention and require manual intervention. This is called “automated with manual approval”—activity is automatic once approval is given. When hardware is on “cold standby,” failover must be performed manually, which invites error.

To date, organizations have primarily relied on BGP (Border Gateway Protocol) as a means for directing traffic over multiple Internet links. BGP was designed for IP routers to direct how packets traversed along the Internet from point A to point B.

While BGP is a core technology for routing, implementing multi-homing using BGP is extremely difficult.

Specifically:

- BGP is costly and complex to deploy - BGP requires hard-to-find network expertise and usually requires costly high-end routers. Organizations often don't have the personnel capable or available for setting up and maintaining a BGP solution. Multi-homing via BGP also requires designated address blocks and an ASN (Automated System Number). In many parts of the world, these numbers are only available to large ISPs, eliminating BGP as an option for all but the largest of organizations.
- BGP introduces latency and performance issues - With BGP, gateway hosts exchange routing information based on data in a routing table. The routing table consists of a list of known routers, the router's known addresses, and each router's path cost metric. Changes to routing tables and the time required for "convergence" of information synchronized between routers can lead to 30 minute delays for changing the direction traffic is sent. BGP, as a result, isn't very nimble - think of a train switch that has to be manually pulled in order to "change the track".
- BGP requires ISP cooperation and maintenance - Often, providers are not willing to set up a "peering" agreement between routers because of significant performance impact to their network. When problems do occur with the configuration, organizations are forced to wait while the ISPs try to determine the cause of and responsibility for the problem.
- BGP provides inferior traffic management capabilities - BGP not only provides a binary traffic decision, forcing organizations to use one primary link for any set of traffic, it also is missing significant capabilities to direct traffic to the best link. BGP has no real control to identify how traffic should be dynamically routed based on line saturation, performance or cost.

For more information, contact Ecessa at 800.669.6242 or visit www.ecessa.com.