CoreXchange

WHITE PAPER

# Increasing the Reliability, Performance, and Value of Internet Connectivity Using Multi-homing

**Maximizing the Internet in a multi-tier environment**

1950 N. Stemmons Freeway | Suite 4006 | Dallas, Texas 75207 | 877-271-7737 | www.corexchange.com

## Introduction

Since its inception, the Internet has been designed not as a single path from one site to another but as a collection of links with routing nodes to decide which link is used to reach a particular destination. As multiple providers built their own backbones, each destination has more and more paths or routes that can reach it and the routing decision is more complex.

Today's Internet consists of dozens of global networks with thousands of smaller providers. Reaching any particular site from any origin can involve hundreds of possible paths and each hop reevaluates the choice of where to go next. An often-misunderstood aspect of Internet routing is the policy or way a path is chosen for your packet. The path is only under your control while it is within your network. Once a packet is delivered to a provider, it is their decision where it goes next. In this way, the modern Internet is really a shared desire to reach each other, with no real guarantee that any one site can be reached.

BGP4 (Border Gateway Protocol) is the core routing protocol of the Internet. It makes routing decisions based on path and network policies.

BGP4 does not use metrics like bandwidth, congestion or latency in its routing decisions

Over time, methods evolved to communicate these routing policies. In the modern Internet, BGP4 is used to convey both what destinations you have on your network, and how you'd liked them reached. It is important to note that this communication is still a request, and it is up to each network provider to decide whether to use the information or not. While this system is very good at allowing new networks to be connected and share paths and policies, it does not really provide for assurances of reachability, latency, or bandwidth.

Once any one collection of paths has been chosen and that first packet makes it to the destination site, the entire process repeats for the response. Another interesting side effect of the current routing architecture is that the set of decisions that led to a particular path being chosen on the way to a site can—and often does—differ for the return path. In other words, the route your traffic takes to a site is often different than the one used to return. This asymmetry is another contributor to the complexity of reaching any one site and adds another layer to the difficulty in assuring that sufficient performance is available for the communication. Therefore, for each destination there is a matrix of outbound and return paths—and this is for any single provider.

## Redundancy and reliability

One of the most common myths of networking is that **redundancy** is equivalent to **reliability**. Actually, no two routers are ever twice as reliable as one. Add to this the fact that "all-or-nothing" failure is rarely the most difficult to deal with, and we can see why reliability can still be an issue. Even with extensive investment in hardware and facilities, any one network link is still vulnerable to either partial or complete failure.

More common today are partial failures. Transient traffic spikes due to re-routing around failed links, DDOS attacks, and human configuration errors can cause partial failures as critical to applications as total hardware failure. As mentioned above, BGP4 may not automatically route around these partial failures since the path still exists, even though performance over that path is degraded.

### DDOS Attack

(Distributed Denial of Service Attack) is an attempt to make a computer system unavailable to its users by over-whelming the computer or its network connections with so much traffic it cannot respond to legitimate requests.

Sophisticated equipment, expertise and relationships with carriers are required to fight these attacks.

## Performance

How does one determine "performance" for an Internet application? Often people confuse simple reachability with performance. Does it ping? In today's applications, we consistently see that simple end-to-end connectivity, while essential, is not enough. If an application is "sluggish" (high latency), then a user will deem the site down in the same way that if it isn't reachable or if the URLs don't resolve. However, simply being quick to return a page is also not enough if the bandwidth is insufficient to send the amount of information in a timely fashion that the application needs. Voice over IP (VoIP), for example, is very sensitive to varying latency and network congestion. Garbled audio and dropped calls cause it to quickly become unusable to end users when there are network performance issues.

In all of these factors, it is important to maintain a minimum of capability in order for the end user's experience to be acceptable. If any of these aspects falls below a critical threshold, the web site or application will be seen to be "down." Given real business and financial realities, however, we cannot pour infinite resources into a particular solution. Cost is inevitably linked to higher performance, and purchasing excess performance leads to excess and unnecessary cost.

So, it becomes obvious that reachability, latency, and bandwidth are essential to Internet performance. But, simple performance is not all that a business requires to succeed. A business needs to accomplish these goals while maintaining the cost component. This is why we speak of value rather than just performance.

## Single Provider Solutions

If a business is using a single Internet provider, it is very difficult for them to control performance and cost. Simply put, a single provider has one price and one network to reach anywhere. This means that in normal operation, there are no choices about how to get to a destination or the performance to that destination. In the event of a performance problem, the traditional approach is to simply open trouble tickets with the network provider. There is no remedy to the problem of cost except negotiating a lower price per megabit. This issue is compounded when changes in your application or customer base change your performance requirements. If suddenly latency to Europe is more important than Asia, then your only solution may be to change providers in order to maintain the performance level that your application needs.

## Multi-provider Solutions

Many early adopters of the Internet quickly realized that the only scalable and maintainable solution to performance and cost issues was to add additional providers to their own network. This way you can get lower cost per megabit and only use the more expensive provider for critical destinations for which the lower cost provider isn't sufficient. You can also increase performance by mixing the best connectivity of two high performance networks to gain better latency or bandwidth than either can provide.

Multi-homing or multi-provider connectivity can solve many of the issues we've looked at so far:

- *Cost* - By adjusting the provider used, multiple prices can be averaged or selected in order to get exactly the performance needed at the minimum cost.
- *Reliability* - Allowing you to route around single provider issues and weather maintenances with little or no impact to your application.
- *Control* - With more providers come more opportunities to affect the return path, and thus half of the performance equation.

### Expense Estimate for Multiple Provider Connections

CAPEX:
- Multiple carrier builds into facility = $20K ea
- Access switches = $30K
- Redundant core routers = $75K

OPEX:
- NOC staff & systems = $20K/month
- Carrier minimum commit = $10K/month ea

But with this technique there is also complexity:

- *Configuration.* In order to leverage any of these benefits you need a configuration to instruct the routing equipment how to use these various providers. The expertise to architect and maintain BGP configurations and links is only one part of the solution. It's also essential to have in-house experts to diagnose and maintain these connections.
- *Hardware.* While competent hardware to maintain full Internet routing tables from multiple providers is dropping in cost each day, it is still expensive to bring in-house and to maintain reliably.
- *Complexity.* Once you have multiple providers for Internet connectivity, any issue becomes one that requires identifying where the issue is located. You must demonstrate this to not just one, but multiple vendors, and the opportunity for misdiagnosis increases.

So while multi-provider configurations do solve many of the issues they also bring complexity, cost, and staffing issues that are daunting even to the largest enterprises.

## CoreXchange Solution

The CoreXchange solution allows a business to outsource the complexity and cost needed to integrate, support, and maintain multiple providers. By bringing customers a single link with the capabilities of multiple providers

integrated into a simple network connection, CoreXchange solves their Internet problem. This solution is a fully managed service, eliminating the capital and operating costs typically associated with multi-provider connectivity. Recognizing that simple multi-provider connectivity still cannot solve all the value needs of every network, CoreXchange goes one step further by providing many multi-provider solutions. Therefore, if an application needs high network performance, the network can be tuned to just as much performance as is needed. If, for example, an application needs mostly bulk throughput, a customer can select a set of providers to bring highest value. Even with all of these options, it is a daunting task for many businesses to select exactly the right network for any one application. CoreXchange addresses this challenge by bringing decades of internetworking experience to bear on the value and performance equations in order to tailor the network to a customer's needs by monitoring and advising the customer continuously on network performance and options.