The Ultimate Online Experience™

INTERNAP

Economics of Multi-Homing and Premise-Based Optimization

Multi-homing, with its assurance of network redundancy, is a critical first step towards ensuring peak performance and 100% availability for your critical Internet applications. Intelligent Route Control allows companies to take advantage of the path diversity that multi-homing provides, to improve network performance and reduce overall network costs. This white paper outlines the initial and ongoing equipment and bandwidth expenditures associated with multi-homing using the Border Gateway Protocol (BGP), and the policy mechanisms that can be employed by intelligent route-control solutions to help leading-edge enterprises achieve its network performance and cost objectives.

An Internap White Paper

Table of Contents

Executive Summary

Purchasing multiple links to the Internet from different service providers (termed “multi-homing”) is a growing trend. Companies initially do so in order to build redundant infrastructures that enable 100% network uptime. However, Intelligent Route Control allows companies to take advantage of the path diversity that multi-homing provides, to improve network performance and reduce overall network costs. Multi-homing, when combined with intelligent route control, promises to finally enable the carrier-class Internet, while optimizing bandwidth costs and providing a cost-effective alternative to expensive private networks.

This white paper outlines the initial and ongoing equipment and bandwidth expenditures associated with multi-homing using the Border Gateway Protocol (BGP), and then discusses the policy mechanisms that can be employed by intelligent route-control solutions to help leading-edge enterprises achieve its network performance and cost objectives.

Introduction

If you control network routing, you control a fundamental piece of your business. The performance of your VPN, the reliability of your web applications and the amount of money you spend on bandwidth are all functions of how your network traffic is routed.

Some would argue that routing is dangerous and complicated – best left to the experts. Certainly, many service providers want you to think so. Prevailing wisdom from traditional service providers is that businesses are unsophisticated, and that providing enterprise network engineers with control over routing is irresponsible, even dangerous to the delicate fabric of the Internet. Of course, this view is also colored by these providers' desire to maintain maximum control over the network and to slow its inevitable commoditization.

Internap believes in intelligence at the Internet edge and sells solutions to enterprises that give them the right tools to control their traffic at the edge as it traverses the Internet. The underlying network is a fundamental part of the business, and that it is too important to leave to blind faith and the best efforts of your service providers.

The Cost of Business

Internet routing determines cost. If you are like the vast majority of businesses, your IP services are billed on the maximum amount of capacity you used during the month over a particular link or links. In measuring this peak capacity used, providers monitor your usage throughout the billing period, usually by taking a measurement of inbound and outbound bandwidth every 5 minutes. The top 5% of usage samples are discarded and the billing is based on the peak of the remaining 95% of the samples. This is referred to as 95th percentile billing.

It is also common to have a minimum commitment level, which sets the floor on how much you would pay per month. In addition, in some cases, customers have a tiered pricing scale, where the unit cost declines in a step function as usage grows. Unless you have a capped-bandwidth tiered scheme or a "full pipe"

FOUR MAIN BILLING MECHANISMS CURRENTLY USED IN THE MARKETPLACE FOR IP TRANSIT SERVICES:

- ***Raw Burstable:*** Customers are charged for the "peak-percentile" (usually peak 95th percentile) bandwidth usage over the billing period at a specific cost/Mbps. In calculating "peak 95th percentile" usage, providers monitor bandwidth usage throughout the billing period (usually taking a measurement of inbound and outbound bandwidth every five (5) minutes and using the larger of the two). And at the end of the billing period, the top 5% of usage samples are discarded, and the peak of the remaining 95% of samples is used to bill the customer. Note that this method of billing based on peak usage means that customers are almost always charged for more bandwidth than they use on average; once a customer peaks to a certain level (for more than 5% of the time), they are billed for that peak.
- ***Burstable with a Minimum Commitment:*** Same as raw burstable, except that the customer has to commit to a minimum amount of usage (usually billed at a flat rate per month), and any amount of "peak-percentile" usage over that minimum commitment amount is billed out at a specific cost/Mbps. As would be expected, the cost/Mbps charged for "committed" bandwidth is considerably lower than the cost/Mbps charged for bandwidth exceeding the minimum commitment, and the greater the bandwidth commitment, the lower the rate.
- ***Tiered:*** In a tiered billing scheme, the customer is charged a fixed rate per month (regardless of the amount of bandwidth actually used), but their usage is capped at a certain level (i.e., the customer cannot consume more than a fixed amount of bandwidth). This scheme offers a predictable and consistent monthly bandwidth spend.
- ***"Full-pipe":*** A specific type of tiered billing scheme, where the customer is charged a fixed rate per month for use of the entire capacity of a given link, regardless of the amount of bandwidth actually used.

scheme from your provider (which is not a cost-effective method to buy bandwidth, and is accordingly not very common), what you spend is determined by how much traffic you send to your provider. And what you send is a function of how traffic is routed.

While bandwidth prices have dropped in recent years, most experts expect that pricing will stabilize over the next year or so. More importantly, however, they forecast that major pricing differences will still exist between providers. The major Tier 1 players will still charge a premium compared to the price charged by pure commodity bit pushers.

Multi-Homing And BGP

If cost were the only criteria you cared about when selecting a provider, then sending all your traffic to a single low-cost provider would probably be the best strategy for you. However, cost is usually one of many considerations. In a recent survey of Fortune 1000 companies conducted by Morgan Stanley, reliability, consistency and performance were all criteria cited as important when choosing a service provider.

Criteria	Avg.	1st	2nd	3rd	4th	Rank
Quality of Service	1.5	71%	17%	6%	5%	1
Price	2.8	17%	32%	23%	20%	2
Network Breadth	4.3	3%	14%	21%	17%	3
Deployment Time	4.8	5%	3%	20%	18%	4T
Technology	4.8	3%	17%	9%	8%	4T
Products and Services	4.8	3%	12%	13%	17%	4T

Source: Morgan Stanley Survey of Fortune 1000 buyers

Although quality was cited as the most important factor for most respondents in the survey, choosing to send all of your traffic over one high-quality provider will not solve the performance problem – no one provider can provide 100% reliable, consistent performance on the Internet. Why? Traffic stays on-net with a single provider less than five percent of the time. On average, traffic passes through more than three networks on its way from origin to destination. Therefore, regardless of how good your provider is, there are almost always other networks that must be traversed on the way from your application to the end-user that will affect the performance of your traffic.

Purchasing transits from multiple service providers or “multi-homing” is the critical first step to controlling the network. A single pipe where all traffic goes in and out to a single upstream provider means that there is nothing you can do to influence how traffic behaves; nor can you guarantee that this Internet link will always be up. Connecting to multiple ISPs allows you to maintain 100% availability, and in many cases, improve performance by getting your application closer to the users who need it.

If you operate your own data centers or have any business-critical applications running over the Internet, it is very likely you are already multi-homed. Two methods used to multi-home are “static routes” and the routing protocol called the Border Gateway Protocol (BGP). With the

former, companies implement “hard-wired” multi-homing using static routes directed to multiple providers. With the latter, businesses multi-home dynamically by establishing peering sessions with multiple upstream providers using BGP.

BGP is the mechanism through which autonomous networks exchange reachability information with each other. If you are running BGP with your providers, BGP announces how to reach various IP address blocks on your network to each BGP-speaking neighbor or peer. This means that if you are running BGP, you are responsible for your own routing and announcements. Your ISPs learn about your network from the prefixes you advertise to them, and this influences how inbound traffic is directed to your application from users located on other networks. Similarly, you learn about how to reach your destination users from the routes you learn from your ISPs.

Even though the entire Internet is reachable from all your ISPs, your edge routers still need to decide which service provider to use. BGP governs how this decision is made. This is an important decision, since most applications have asymmetric traffic flows with the bulk of bandwidth consumed (e.g., 80-90%) on the outbound path. This intuitively makes sense: your users/customers send small messages requesting data with the bulk of content flowing from your application to the end-user. Therefore, outbound (egress) path selection has significant implications on your cost. However, BGP is a deterministic protocol, in that it applies pre-defined rules to make routing decisions. These rules are based on a hierarchy of decision variables, some of which can be manually influenced to meet your objectives, but only in a coarse-grained, periodic manner – and without any awareness of actual traffic performance or the associated bandwidth cost.

In order to take control of your routing, you must be able to control BGP in an automated fashion and direct traffic over links that meet your business needs, including cost and performance criteria. Before examining how you can get this control, let us look at the costs of multi-homing.

The Cost of Multi-Homing

If you are not currently multi-homed, the out-of-pocket costs to migrate from single-homed to a multi-homed network include the following one-time and ongoing costs:

1. *Layer 2 access to your additional ISPs.* If you are connecting to another service provider, you need a physical connection to this other network. The costs vary by access method and location, but some representative costs are:

    - For a short-distance T1, the monthly recurring costs are approximately $475 per month, with a one-time installation charge of about $500.

- If you are in a third-party data center and want access to another resident carrier, a 100 Mbps Ethernet cross connect will be on the order of $200 per month, with a $250 installation charge.

2. *Additional router line card.* Each physical link must terminate somewhere, so an additional line card is often required for your router depending on the router model and the link type. For example:

   - For a Cisco 2800 (series), a T1 line card is approximately $700 and a 100 Mbps FE port adapter is approximately $800.
   - For a Cisco 7300 (series), a T1 line card is about $3,000 and a single port 100 Mbps FE line card is about $1200.*

   * Please note that the provided costs are examples only using U.S. prices and Cisco equipment. Costs outside the U.S. will vary widely, as would equivalent equipment from other vendors.

If you plan to multi-home using BGP (rather than static routing), additional costs include BGP router upgrades, configuration and setup fees:

3. *Router upgrades.* To receive full routes, your routers participating in BGP need a minimum of 512 Mbps of RAM. You may also need a code upgrade to accommodate BGP, although a Cisco router with 10.3 or later IOS will do.

   - Most businesses that utilize or are considering a multi-homed topology will have a router chassis capable of operating in a BGP environment. If not, then a device such as a Cisco 7200VXR(series) with NPE-G2 Route Engine and Fast Ethernet line card(s) would be more than sufficient for most needs. This would cost approximately $15,000.

4. *Autonomous System Numbers.* Autonomous Systems (ASs) are IP networks that adhere to a single routing policy. Each network must have a globally unique Autonomous System Number (ASN) in order to exchange exterior routing information using BGP. In the U.S., AS numbers are controlled and issued by the American Registry of Internet Numbers (ARIN); there is a one-time registration fee of $500 for each AS number assigned. Usually only one is required.

An interesting note: the number of Autonomous Systems has been growing at more than 50% per year – and this growth rate is accelerating, suggesting that BGP-based multi-homing by the large enterprise is becoming a common network architecture.

5. *IP address block registration.* While you may obtain IP address space from your ISPs, you may also request your own address blocks from ARIN. The fee for addresses is based on the amount of space assigned. For example, being assigned a /20 is a $2500 one-time cost.

It is possible to move from a single-homed environment to a multi-homed topology for around $4,000 -$5,000 up-front with additional monthly charges of $500 per link, assuming your edge routers are relatively new and well provisioned.

What do you get for this extra cost? First, you get physical link diversity with two or more ISPs and the benefit of fast fail-over, thanks to BGP recognizing that a link is down or some upstream prefix is not reachable by the current route and automatically rerouting your traffic. You are also likely to realize some performance gain (assuming that you selected your ISPs wisely), since the number of hops required reaching a given destination often varies between different providers.

However, there are additional, hidden bandwidth costs that you pay when you migrate from a single-homed environment to a multi-homed environment.

Inefficient peak utilization. Since the distribution of peak traffic loads over multiple links tend not to happen at exactly the same time (in-phase), inefficient 95% percentile peaks occur across multiple providers. These multiple peaks cause you to pay for more bandwidth than what you would pay if you used only one link. It is not uncommon for customers to pay for 2X the amount of bandwidth than what they otherwise would pay if they were sending the same amount of traffic over only one provider.

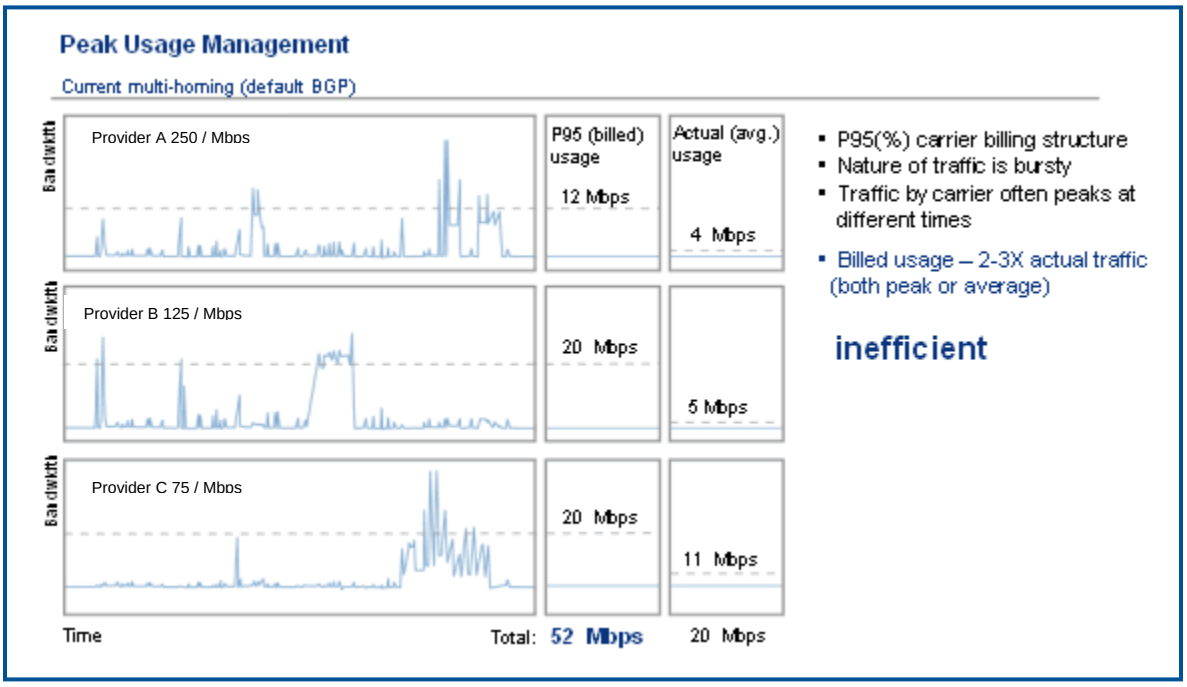

6. *Highest-cost routing.* BGP directs traffic according to its hierarchy of variables. While there are some elements that are under your control, what tends to happen in practice is that the decision variable that ends up being used to select a route for any particular egress flow from your application is based on least number of Autonomous System

(AS) hops. The providers that tend to have the most extensive reach and peering relationships tend to have the fewest AS hops to a particular address. These also are the larger, more established – and most expensive providers. Therefore, while a lower-cost alternative may exist, BGP will force traffic over the more expensive link, bypassing the adequately performing, cheaper link.

In summary, while there are major benefits associated with being multi-homed, there are clearly costs associated with it. However, the major cost item is not the physical connectivity to additional networks – it is the bandwidth costs, which are a function of how traffic is routed amongst your ISPs. Therefore, if you can control how your traffic is routed, you cannot only control its performance, but you can control your spending on bandwidth.

Saving Money with Premise-Based Route Optimization

The Internap® premise-based optimization solution saves our multi-homed customers money, not just soft costs like the cost of downtime, lost revenue, employee productivity, and so on, but real out-of-pocket cash, in some instances up to 40%. Internap's position on intelligent route control and multi-homing is better performance is good, but better performance with cost savings is better.

Internap offers a solution that can optimize significant incremental cost, while simultaneously delivering consistent, predictable performance. In fact, Internap customers, who are some of the most advanced and sophisticated network buyers in the world, have experienced 15% to 40% bandwidth optimization when deploying the Internap premise-based route-control solution.

Four cost-optimization mechanisms can be used to save money. The degree to which they apply to your situation will vary depending on your traffic distribution profile, your specific contracts with your service providers and the level of performance gain you seek to realize out of your current network links.

1. *Routing to the least cost provider (subject to a performance constraint)*
   The simplest cost optimization technique is to employ least cost provider and performance awareness into the route decision-making process. A material price differential does exist between service providers. In the case of one of our Fortune 100 customers, the difference between the most expensive ISP and the least is 240%. Therefore, it is in this customer's interest to use the lower cost provider as much as possible and the most expensive provider as infrequently as possible. In this customer's case we were able to redistribute their traffic profile as follows:

	Provider Relative Cost	Before (BGP)	After (Internap)
ISP 1	100%	37%	78%
ISP 2	160%	17%	10%
ISP 3	240%	46%	12%

This reduction in costs is in conjunction with improving their performance. Least-cost routing alone makes no sense. If cost is your only consideration, sending all of your traffic over the lowest cost provider will meet your needs. This customer-defined latency and packet-loss thresholds that were enforced across its three service providers. And, in almost 80% of the cases, the inexpensive provider met the performance requirements of this particular application. Only when it did not, did the Internap solution re-route to a more expensive provider. This route activity, which delivered consistent performance while employing a kind of cost arbitrage, yielded our customer monthly bandwidth savings of more than 35%.

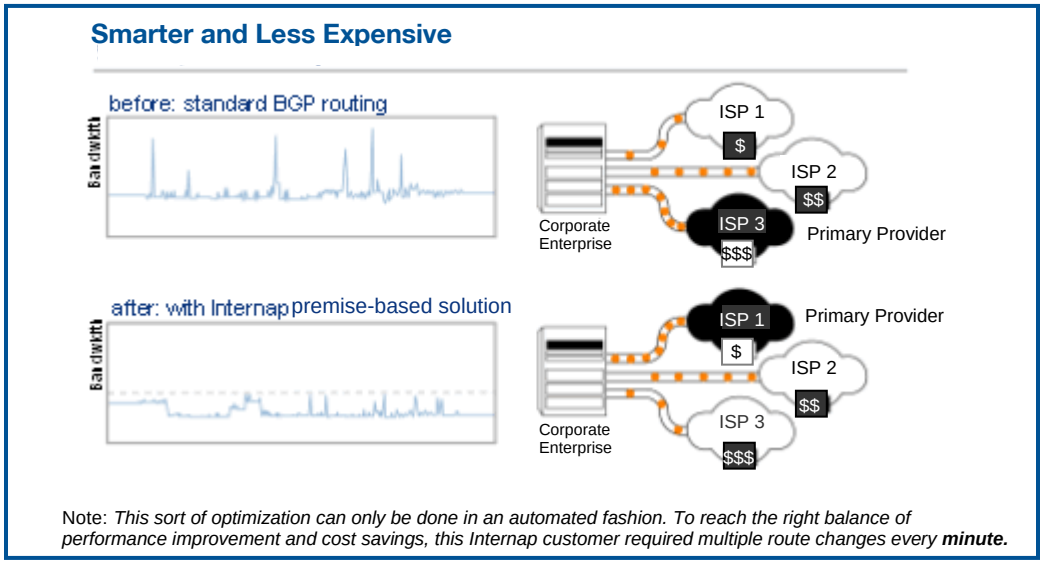

Note: *This sort of optimization can only be done in an automated fashion. To reach the right balance of performance improvement and cost savings, this Internap customer required multiple route changes every* ***minute.***

2. *Minimum commitment and peak usage balancing (subject to performance constraints)* If you just worked on the basis of "use my least-cost provider and only move traffic off of it if it exceeds performance thresholds X, Y and Z," controlling costs would be a simple exercise. However, most customer contracts require an intelligent view of marginal cost. For example, if you have minimum commitments with multiple providers (you probably do), then this is "paid for" bandwidth regardless of whether you use it. Therefore, what you want from a routing perspective is a policy that states: "use up all of my minimum commitments before bursting above them – as long as the providers meet my performance criteria."

   Even this is a bit simplistic. It is likely that you are charged based on some peak usage scheme (e.g. 95th percentile billing). Therefore, not only do you have fixed minimum commitments to take into account, once you have peaked above this level, all

subsequent traffic under this peak is effectively free. Unlike minimum commitments, the peak during a given period requires keeping track of a running calculation of your maximum level of traffic over every transit link over a period of time – and recognizing whether a route change for a specific flow will cause you to create a new "peak."

Without an intelligent route-control system, BGP will create multiple peaks, across your providers. The protocol will make its own decisions and distribute traffic from your application without regards to minimum commitment and peak usage levels. Based on Internap's experience with large customers, if you have a multi-homed networks using BGP without intelligent route control, you are likely paying an excess of 50 – 200% more for bandwidth than you would under optimal routing circumstances.

Done right, least-cost routing is based on marginal cost decision making on a prefix-by-prefix basis. And, this is driven by where you are, in real-time, in your peak-billed usage levels. Again, all are subject to performance constraints.

Through this mechanism, the over-billing of bandwidth can be reduced from, for example 2X, to a more typical level of 1.2X. Instead of paying for 200 Mbps per month, you pay for 120 Mbps – and at a lower blended rate per Mbps, because lower cost providers turn out to work just fine a large percentage of the time.

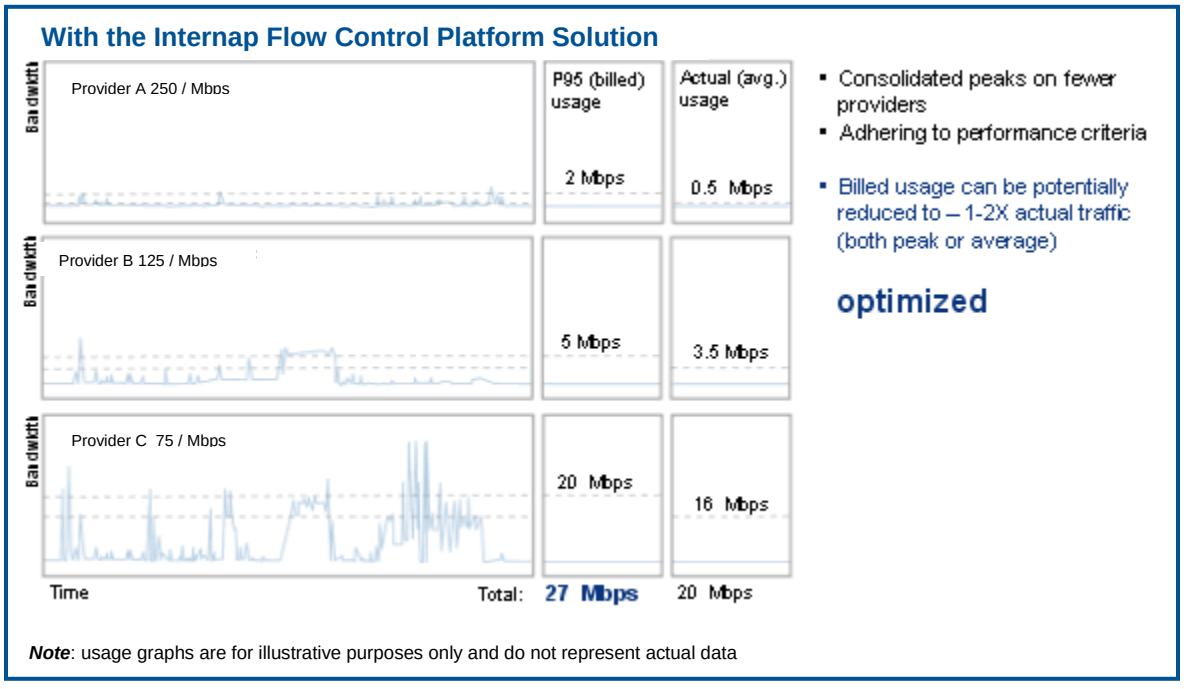

Note: usage graphs are for illustrative purposes only and do not represent actual data

3. *Right-sizing network architecture*
   Lastly, once you control how your routing behaves, you can optimally structure your relationships with service providers. If you know you can direct the majority of traffic

and manage peak levels with a low cost provider, you can structure your minimum commitments to them in order to further drive down cost per Mbps. For the other, more expensive providers – since it is likely that you will only use them sparingly, you can restructure your contracts accordingly. This might take the form of lower commitments, but even with higher per Mbps prices you are likely to save money because you send dramatically less traffic over their network.

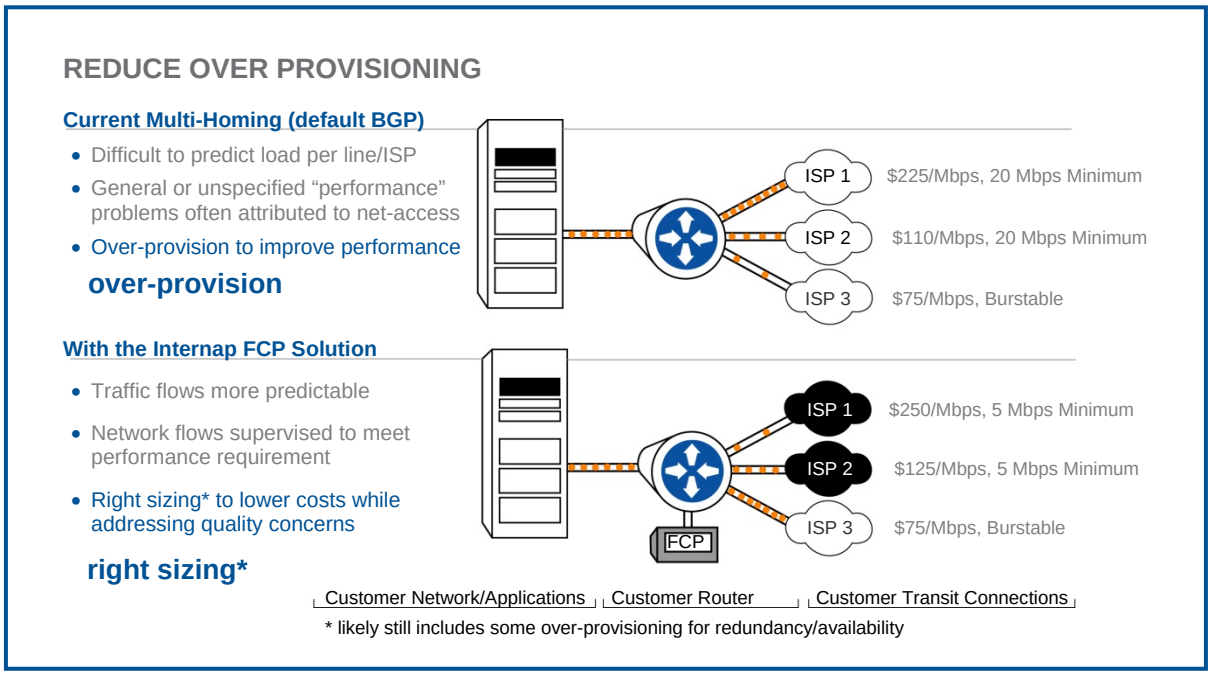

4. *An even bigger opportunity – IP VPNs*
Customers with significant internal "corporate" networking requirements can exploit intelligent route-control technologies to go beyond optimizing the cost and performance of their web applications, to migrating traffic from legacy private network links onto Internet based IP VPNs.

In our experience, there have been two problems that have kept data traffic on legacy networks. The first is security: leased lines, ATM and Frame Relay have security as their hallmark, and were traditionally the only choice to ensure that business-critical enterprise traffic remained secure. However, with advances in encryption, and intrusion detection, the security concerns associated with the Internet have started to fade. The Internet today is sufficiently secure for almost any application.

A more serious problem has been the unpredictable end-to-end performance of the Internet. IP still has inherent limitations when compared to Frame relay and ATM. Neither end-to-end bandwidth prioritization and reservation nor end-to-end performance and availability guarantees are available.

The Internap Flow Control Platform™ premise-based solution eliminates this end-to-end performance barrier. Our route-control technology delivers *consistent* performance end-to-end, to any destination on the Internet. Suddenly, your legacy network's expensive committed information rate can be substituted with a much cheaper "optimized" IP connection that can deliver end-to-end performance within defined performance boundaries. The decision is made easier, given that a partial-mesh Frame Relay network's per-port cost 1.5 to 2.5X the cost of ISP connectivity, and private line networks are 4 to 6X the cost.

In addition to saving 50 to 200+% off your private networking costs, additional benefits include simpler administration and the ability to exploit the Internet's ubiquity to connect business partners, customers, remote office locations, work-from-home employees and foreign locations to your network. Furthermore, you can also gain economies of scale in bandwidth, equipment and labor costs by consolidating your private network traffic (and if you are really aggressive, your internal voice traffic) with your Web/IP network.

Summary

	Costs	Benefits
Step One: From single to multi-homed	• Layer 2 connectivity to ISP's • Additional router line cards • ASN and IP address space • BGP caused layer 3 inefficiencies	• High availability and fast failover • Better average traffic performance • Vendor leverage
Step Two: From basic multi-home to active route control	• Route-control system	• Private network quality end-to-end • LCR and optimal traffic redistribution • Efficient peak utilization • Right-sized networks and optimal contracts • SLA and vendor management • Reduced manual effort
Step Three: From legacy network to route-control enabled IP VPN	• IP VPN security infrastructure/CPE • IP VPN provisioning and management apps	• Lower per-port cost • Extend network reachability • Ease of administration • Economics of scale in converged network

Multi-homing, with its assurance of network redundancy, is a critical first step towards ensuring peak performance and 100% availability for your critical Internet applications. However, only premise-based intelligent route-control options like the Flow Control Platform solution from Internap can provide the necessary tools to unlock multi-homing's additional benefits: reliable, predictable Internet performance, improved bandwidth cost control and a ubiquitous cost-effective alternative to expensive private networks. Companies that deploy the Internap solution can now fully control its BGP routing behavior without extensive training or manual effort and can gain the performance and usage visibility it needs to tailor its traffic routing policies to meet business needs and effectively manage its ISP relationships. The carrier-class Internet is finally a reality.