

Topology Aware Route Control: Understanding True Path Diversity

Many route-control products on the market today utilize and make control decisions on the routes as advertised in the current Internet routing table. Unfortunately, these route-control products do not take into consideration the vast geographic distances between adjacent networks of a large address block. Today, a solution exists that delivers uncompromising performance optimization of even large address allocations. This white paper explores the Internap approach to its IntelliPoint Routing™ technology as well as the concepts and motivation behind this unique technology.

An Internap White Paper

©2008 Internap Network Services Corporation. All rights reserved.



Table of Contents

Executive Summary.....	3
Introduction.....	4
Convergence Point Analysis	6
True Diversity	8
Last Mile Effects	9
Intelligent Advertisements	10
Summary	11



Executive Summary

There are not many route-control products on the market today that utilize and make control decisions on the routes as advertised in the current Internet routing table. These routes are often large allocations of address space meant to keep the inter-provider routing table small. This aggregation of routes is an important criterion when the routing tables are communicated among large Internet Service Providers (ISPs), as is the primary goal of the Border Gateway Protocol (BGP).

Unfortunately, these route-control products do not take into consideration the vast geographic distances between adjacent networks of a large address block. At best, these products force customers to make a route-control decision that degrades some portions of the address space “for the greater good.” At worst, when the product is unaware of this address distribution, the customer faces control decisions that actually introduce larger problems than are resolved.

Today, a solution exists that delivers uncompromising performance optimization of even large address allocations. This solution is the Internap® IntelliPoint Routing™ technology. The IntelliPoint Routing technology leverages network topology information to pinpoint the precise location of middle-mile performance problems and target decisions that only affect the address space experiencing the performance problems. This white paper explores the Internap approach to the IntelliPoint Routing technology as well as the concepts and motivation behind this unique technology.



Introduction

With the introduction of classless inter-domain routing (CIDR), the routing table is built using a variety of network sizes. To illustrate this, consider a network that contains many IP addresses, like a corporate LAN. This network (or a group of such networks) is listed in the routing table as a network prefix. A prefix is just a 32 bit IP address that has an associated netmask indicating how many of the leading bits are significant.

BGP4 is the first version of the protocol deployed to handle the variable length prefixes introduced with CIDR. With BGP4, a prefix no longer has to fall on a byte boundary (Class A, B or C space), but can be one of 32 different sizes depending on the length of the network mask. The common notation for a prefix is “address/netmask” (e.g. 12.0.0.0/8). The 12.0.0.0 is the address and the “/8” indicates that only the first 8 bits of that address are significant, in this case, just the “12”. The fewer the bits in the netmask, the more IP addresses in the prefix. For example, a /16 prefix has 256 times the address space of a /24 prefix given the 8 bit difference in the netmask.

Many route-control products on the market today utilize and make control decisions on the routes as advertised in the BGP routing table. These routes are often large allocations of address space meant to keep the inter-provider routing table small. Unfortunately, these route-control products do not take into consideration the vast geographic distances between adjacent networks of a large address block.

The problem is easily understood when the address allocations in a single aggregated route fall into widely different locations as shown in Figure 1 below. In this figure, one of the address allocations occurs on the east coast and another on the west coast. It is uncommon that a single route decision for the /16 can effectively optimize every address in the block. At best the customer will face a performance degradation at some destinations “for the common good” of all other destinations. At worst, the product will fail to recognize this geographic diversity. And, on fixing small problems for a particular destination will introduce greater problems for a larger set of destinations. If this effect is not recognized, the product may then introduce performance route flapping for a large address block. This effect is demonstrated in Figure 2 on the following page and was seen during a competitive product evaluation at a large enterprise customer. In this figure, control decisions on just 20 routes in the routing table move large volumes of traffic between two providers. Clearly, any product that delivers such heavy-handed advertisements can be quite disruptive to enterprise networks.




Figure 1 – Address allocations in a single aggregated route that fall into widely different locations



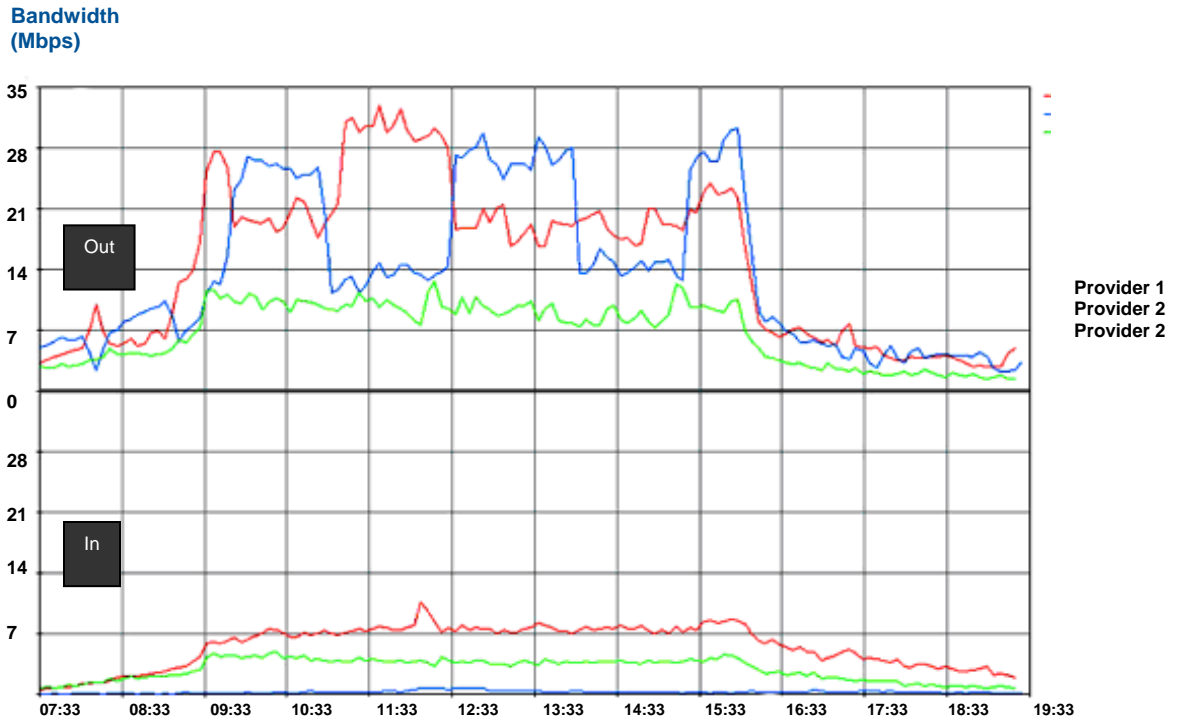


Figure 2 – Effect during a competitive product evaluation at a large enterprise customer

Figure 3 shows that the majority of Internet address space falls in prefix lengths shorter than /17.

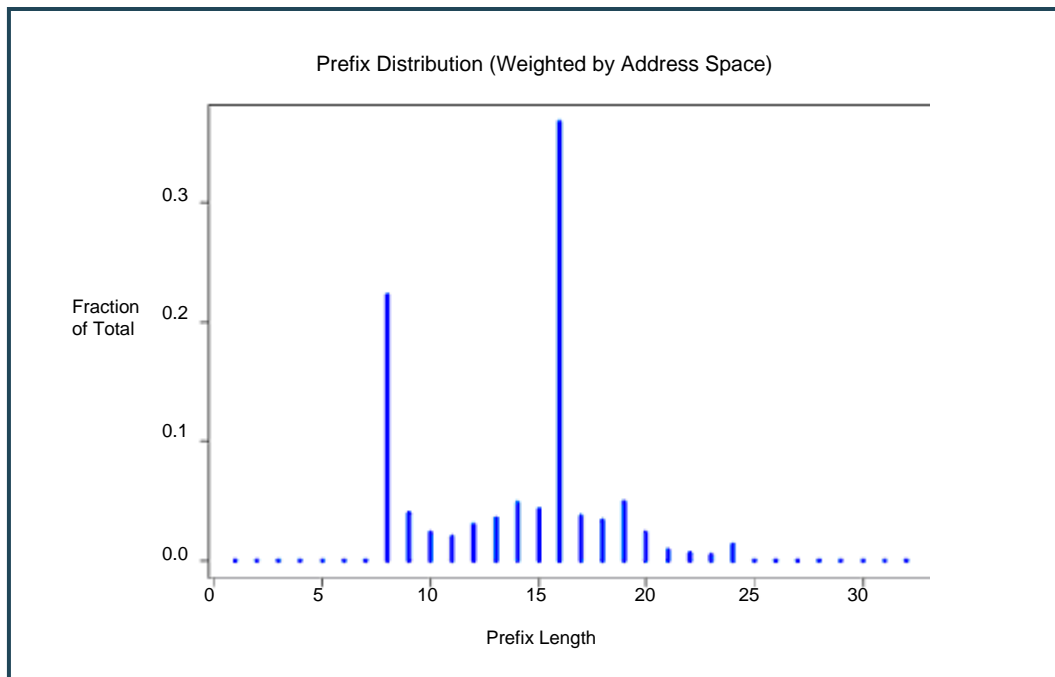


Figure 3 – Majority of Internet address space falling in prefix lengths shorter than /17



These statistics are easily illustrated simply by using traceroute within some common address blocks. Observing the path diversity to a few specific /24s in any /16 will show that the address space is geographically diverse (from both latency and DNS names). Effective route-control solutions must be aware of this route diversity in order to make truly intelligent decisions.

Today, a solution exists that solves not only the aforementioned problems but also delivers greater economies of scale and more intelligent decisions by leveraging the information gathered about the topology of the middle-mile. This solution is the Internap IntelliPoint Routing™ technology. The IntelliPoint Routing technology delivers more effective and proactive route decisions to a larger set of destinations while minimizing the amount of measurement required to make those decisions. While alternative products deliver at best a compromise and at worst large-scale performance degradation, the Internap Flow Control Platform™ (FCP) solution delivers accurate and appropriate changes to the routing table.

This white paper examines the IntelliPoint Routing technology and illustrates how this technology delivers:

- Significant gains in measurement capability
- An understanding of provider diversity
- Differentiation of middle-mile and last-mile effects in all decisions
- Intelligent advertisements

Convergence Point Analysis

The IntelliPoint Routing technology leverages convergence point analysis to deliver the aforementioned benefits. Convergence point analysis identifies unique network points at which all available paths converge towards a destination. This represents the point in the network where path diversity effectively falls to zero. Analyzing the full path information already being collected in the FCP solution today enables the discovery of this convergence point.

Figure 4 on the following page shows the convergence point analysis of three paths to a given destination. Two of the paths converge at point A while all paths converge at point B. All remaining hops beyond Point B to the destination are the same for all candidate paths. If you want to measure the portion of the middle-mile that is diverse and understand the real difference between the available providers in the middle-mile, point B is the logical point to measure. Measuring points beyond point B includes common hops for all providers, which can cloud the comparisons. Measuring points before B does not consider all of the diverse hops.



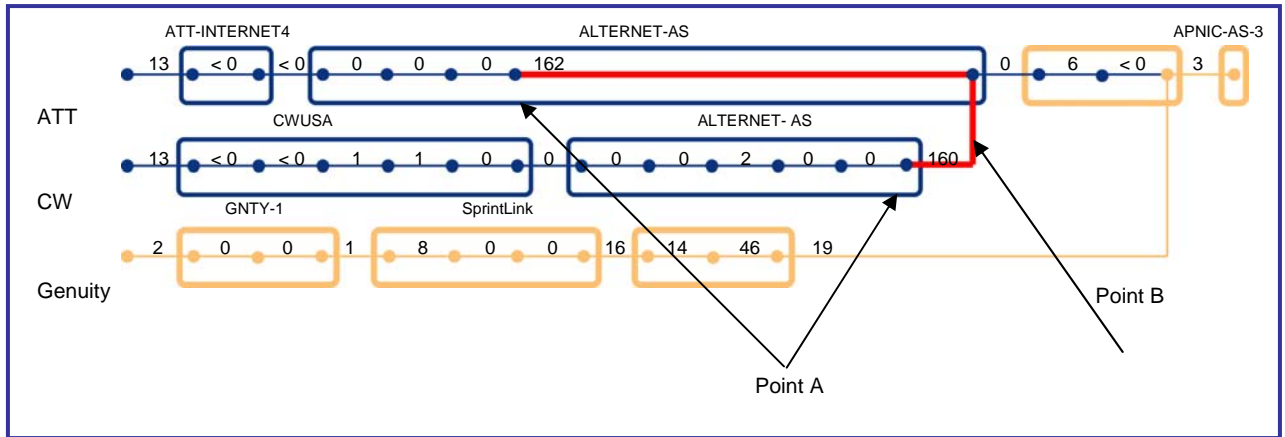


Figure 4 – Convergence point analysis of three paths to a given destination

Figure 5 shows another convergence point but expands the set of addresses observed beyond the convergence point. In this diagram, there are several destination networks that share the same convergence point. In understanding all of the networks that are beyond a convergence point, The IntelliPoint Routing technology is able to aggregate a single measurement made to the convergence point and report the results for all of the addresses beyond the convergence point. With intelligent advertisements, which are discussed below, the IntelliPoint Routing technology makes only one route change for all of these networks. This approach delivers unprecedented scaling benefits to a solution that already leads the category in scale today. It also serves to lessen the impact of active measurements on remote security boundaries and local egress bandwidth since lower aggregate probe volume is required to achieve more effective results.




Figure 5 – Convergence point expanding set of addresses observed beyond the convergence point

The scaling benefits of the IntelliPoint Routing technology become readily apparent when considering the network aggregation illustrated for the convergence point of Figure 5 over all convergence points of a typical customer.



Figure 6 shows the number of destination networks behind the top convergence points observed over an hour at an example customer. The effects are fairly uniform with an average of 100 networks behind each convergence IP. Measuring just the convergence point yields a 100 to 1 gain in scaling, that is, a single measurement today can do the work of 100 measurements of the past.

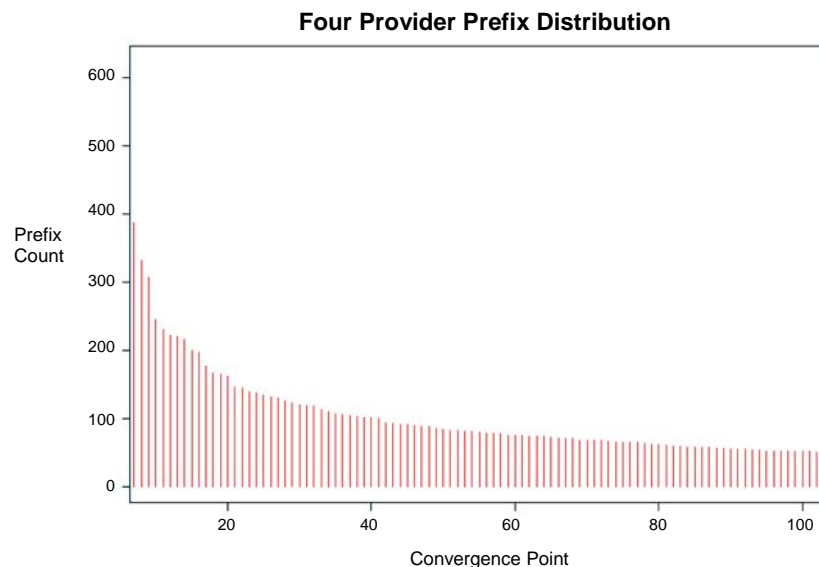


Figure 6 – Number of destination networks behind the top convergence points observed over an hour at an example customer

True Diversity

With its convergence point analysis, the IntelliPoint Routing technology is able to highlight the true diversity being offered by the current transit providers. This helps the engineer select not only the correct number of providers, but also which specific providers are best at getting to the destinations of interest. As will be demonstrated, there may be peering relationships coupled with hot potato routing, which prevent the destinations from having any path diversity.

The data in Figure 6 was collected from a content site that has four service providers. When analyzing the effects of all convergence points, the distribution of networks per convergence point is fairly uniform. When the number of providers is reduced, a peculiar thing starts to happen. Some destinations begin to see little or no path diversity. This can be seen when large numbers of destinations (>1000) converge close to the customer site. This is shown in Figure 7, which graphs the distribution of networks behind convergence points of a customer with only two transits and shows several high-density convergence points. When probing the distance to these high-density convergence points, we learn that they are only a few hops away, leaving the majority of the network path with zero diversity for these destinations.



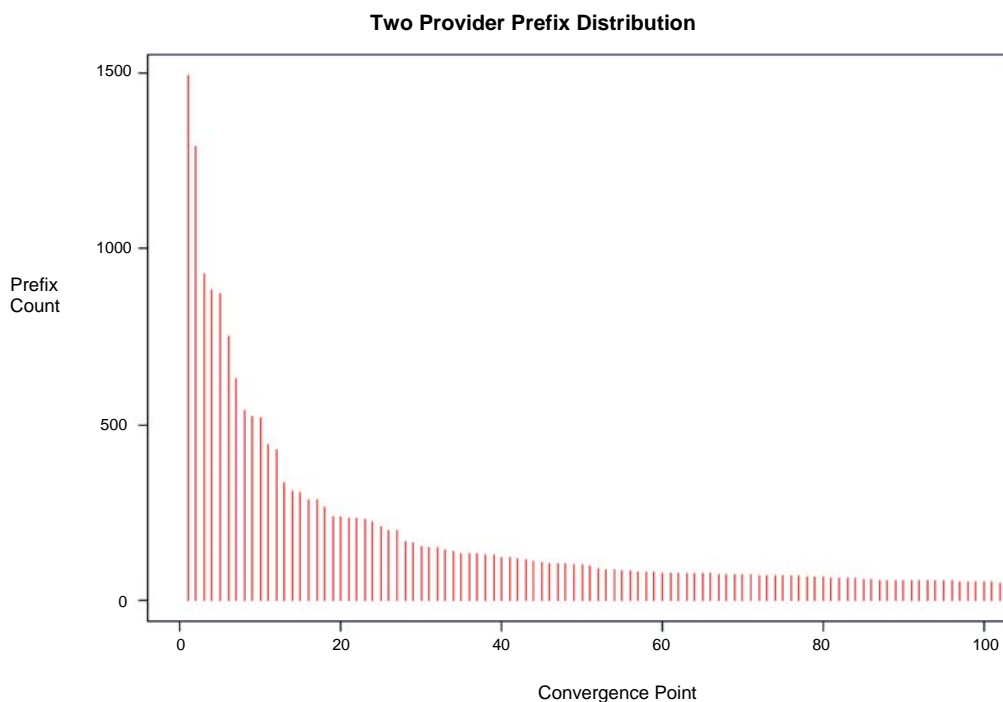


Figure 7 – Distribution of networks behind convergence points of a customer with only two transits and shows several high-density convergence points

In this example, this effect is the result of two forces, a nearby peering point between the two transit providers, and hot potato routing. When Provider A gets traffic destined for Provider B it directs it “hot potato” to the nearby peering point where the path quickly converges with the path via Provider B. The same effect happens for Provider A traffic sent out Provider B. But selecting a different provider pushes the convergence points out substantially, and brings more diversity to all destinations.

To maximize the benefits of route control, the diversity offered through available providers should be analyzed. The convergence point is the most effective way to accomplish this. To help the network engineer perform this analysis, the Flow Control Platform solution offers special representations that reports on the diversity being offered by the current providers. This planning tool empowers the network engineer with the information necessary to make intelligent provider selections. This, in turn, will maximize network diversity; and in doing so, increase network performance.

Last Mile Effects

End-to-end measurements are good at indicating deviations from the norm and detecting performance problems. However, effects on the non-diverse portions of the path can cloud end-to-end provider comparisons. End-to-end route-control techniques such as a web object are often confused by last mile congestion, which in turn can cloud the routing decisions being made among a set of providers.



At times, the convergence point occurs before the portion of the path that is impacted as shown in Figure 8. In these situations, it is not possible for route control to select a path around the problem. But nevertheless, it is important that route control be aware of these situations. This prevents the product from making route changes that attempt resolution of problems it cannot solve.

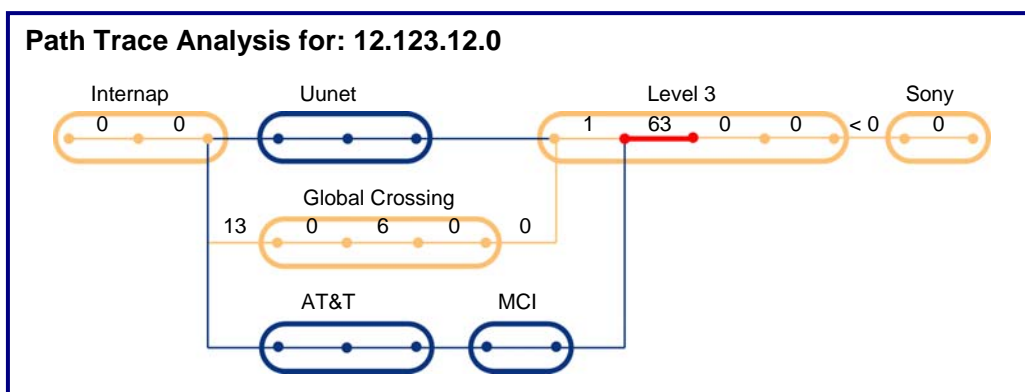


Figure 8 – Convergence point occurs before the portion of the path that is impacted

Unfortunately, unless the product recognizes these situations by measuring only the portion of the path that is relevant (that is up to the convergence point), spurious route changes and churn will occur. Experience has taught us that under congestion, the variance in latency can be very high. One sample may vary by 100s of ms from the next, to the same destination, as part of the same flow. Using the IntelliPoint Routing technology, these effects can be isolated to a last-mile link (or overseas link for which there is no diversity), obviously post-convergence point. While the Passive Flow Analyzer™ tool analysis observes these as performance problems, the IntelliPoint Routing technology quickly recognizes that the flow is already on the best path. If the route-control product assesses candidate paths in an end-to-end manner with a web object, the candidates would appear to differ by the same 100s of ms, packet-by-packet. This will cause some paths to appear better at some times and other paths to appear better at others... all from last-mile effects!

The IntelliPoint Routing technology, by assessing the true path diversity, is able to weed out the problems that cannot be solved without inducing needless change or route flapping to the network.

Intelligent Advertisements

Using convergence point analysis, the IntelliPoint Routing technology is able to announce routes at the appropriate aggregation based on topology. This is accomplished by maintaining and storing the convergence point for each unique network block that the system has observed. When an announcement should be made for a particular precise network block (say a problem observed at a particular destination, e.g. 12.0.1.0/24 of Figure 5), the convergence point is compared to all adjacent network blocks to determine if any aggregation is possible. If two or more adjacent blocks share a common convergence point, the aggregated advertisement is made instead.



For example, if destination networks 12.0.0.0/24, 12.0.1.0/24, 12.0.2.0/24, and 12.0.3.0/24 all share a common convergence point, it is possible to aggregate the address space and announce 12.0.0.0/22 when making a route change for any of these networks. Because the convergence point is the same, there is no danger that one of the networks is geographically distant from the others (avoiding the common problem when using the routes of the BGP routing table). When a problem is observed at one of the destination networks (say 12.0.1.0/24 in the example) the system can proactively move all of the addresses that share the destination's convergence point (in this case, 12.0.0.0/22). This proactive method often resolves problems for many destinations well before they are ever observed. Networks with common convergence points are an effective way to represent and store the amount of address space that can be safely moved in response to any problem. Moving the address space with common convergence points will never introduce problems that occur when moving routes as seen in the routing table.

We have seen from Figure 6 that it is not uncommon for a hundred networks to share a common convergence point. Reacting to a single problem in one of these networks can proactively improve performance for the remaining 99 destinations.

With these intelligent advertisements, the Flow Control Platform solution proactively resolves problems for most addresses well before they are encountered all the while dramatically reducing the number of announcements required to resolve large scale problems in destination networks.

Summary

Using convergence point analysis, the IntelliPoint Routing technology delivers uncompromising resolution of middle-mile network problems, greater economies of scale in both measurement and resolution capacity, while reducing the overall operational overhead of route control with more selective and more intelligent measurement techniques. Leveraging this topology information results in the most intelligent decisions that can be made and delivers a wide array of benefits to the enterprise customer that wants to improve performance, optimize network costs and exercise the utmost control over its network traffic.

