INTERNAP®

The IntelliChoice Model for Both Cost and Performance

When routing Internet traffic, best performance per destination and overall cost optimization objectives have traditionally been in conflict. With most premise-based optimization solutions available today, the user must select either cost or performance objectives for each destination. A solution exists that can both improve performance and lower costs at the same time for all destinations. This white paper examines the evolution of premise-based optimization and details the key technologies within the Internap® Flow Control Platform™ (FCP) solution that enables the IntelliChoice technology, the Passive Flow Analyzer™ tool and simultaneous cost and performance optimization.

An Internap White Paper

Table of Contents

Executive Summary

When routing Internet traffic, best performance per destination and overall cost optimization objectives have traditionally been in conflict. Because there is no single Internet service provider that can provide high performance and lowest cost delivery to every destination on a national or global scale, these objectives are nearly impossible to achieve simultaneously without purchasing links from multiple Internet service providers, “multi-homing” and a premise-based optimization solution.

But optimization solutions which attempt to improve performance and optimize costs for Internet traffic have had similar cost/performance conflicts. These solutions rarely achieve both cost and performance gains for the same destinations, forcing network managers to make costly business tradeoffs. In fact, with most premise-based optimization solutions available today, the user must select either cost or performance objectives for each destination. As such, cost-based route changes do not consider any performance implications and performance-based route changes do not consider cost implications.

A solution exists that can both improve performance and lower costs at the same time for all destinations. This is combined with measurability of your network traffic which all combines to truly understanding your network and managing it effectively. The Internap® Flow Control Platform™ (FCP) solution, with its IntelliChoice™ technology, continuously considers performance and cost implications by applications with every route change for all destinations. This white paper examines the evolution of premise-based optimization and details the key technologies within the FCP solution that enables the IntelliChoice technology, the Passive Flow Analyzer™ tool and simultaneous cost and performance optimization.

The Evolution of Premise-Based Optimization

Premise-based optimization for enterprises began with the extension of the Border Gateway Protocol (BGP) and multi-homing from pure service provider networks to the corporate enterprise. The goal was to provide backup network access in the event of service failure. Yet, even today, the BGP decision process does not consider cost or performance measurements when directing traffic for a multi-homed enterprise. Basing its decision mainly on the shortest number of intermediate provider networks from source to destination, BGP frequently raises network costs and lowers performance needlessly by not considering either variable. Specific premise-based optimization solutions were introduced to overcome these problems by intelligently routing traffic according to customer-defined policies based on optimized cost or performance thresholds.

Second generation premise-based optimization technologies create business policies that allow customers to route traffic based on either cost or performance metrics. These products tend to rely on Cisco Netflow for limited, basic network usage information; service providers are ranked based on the average cost of bandwidth. Traffic flows are apportioned according to either usage information or provider ranking. While such a solution is unlikely to produce the absolute minimum cost for an enterprise, this simple optimization can reduce some cost or provide marginally better performance for network users.

The Internap Flow Control Platform (FCP) solution with the IntelliChoice model represents the third generation of premise-based optimization technology. Its advanced understanding of provider billing contracts and ability to track current usage can produce the absolute minimum cost for a given enterprise and its network traffic. Employing unprecedented scale in real-time processing, the IntelliChoice technology can combine billing data with in-depth information about each traffic flow to maintain optimum performance while also minimizing cost. Utilizing the FCP Passive Flow Analyzer tool, described later, the IntelliChoice model can recommend flow swaps that reroute traffic for performance reasons, even when providers are reaching cost thresholds. New capabilities of the FCP solution allow for these changes to be done on an application level so certain applications always have access to the best performing or least-cost routes The IntelliChoice-driven performance swaps guarantee the absolute minimum cost while maintaining desired performance levels.

This white paper examines the Internap FCP solution with its IntelliChoice technology. We will explore how the FCP solution’s unique, scalable architecture and design enables enterprises to achieve the lowest cost and the best performance for all of their traffic at the same time.

The FCP Architecture

While most of this white paper focuses on the IntelliChoice technology component of the FCP solution, it will be important to understand the overall architecture of the FCP solution as well. The solution is engineered to deliver cost and performance-based route control, for either one variable or both variables simultaneously. It is built around a highly-scalable architecture that provides real-time visibility into the four variables required to optimize cost and performance: flow volume, flow performance, provider usage and provider cost.

Figure 1 below shows the architecture of the FCP solution. One of the major components is the FlowCollector™ module with the built-in Passive Flow Analyzer (PFA) tool. The PFA tool provides the FCP solution with real-time flow performance and volume information that is necessary for both cost and performance decisions. The ability to improve performance and optimize cost at the same time lies in the PFA tool's capability to monitor hundreds of thousands of concurrent flows and over 100,000 simultaneous network prefixes. This clearly enables the FCP solution to catch any performance problem on any of these destinations. It also provides detailed information on a very large number of candidate flows that can be moved to resolve cost- or performance-based violations, as described in this white paper.

No other solution offers the same visibility into all of your IP traffic. This unique technological advantage enables true concurrent cost optimization and performance-based control.

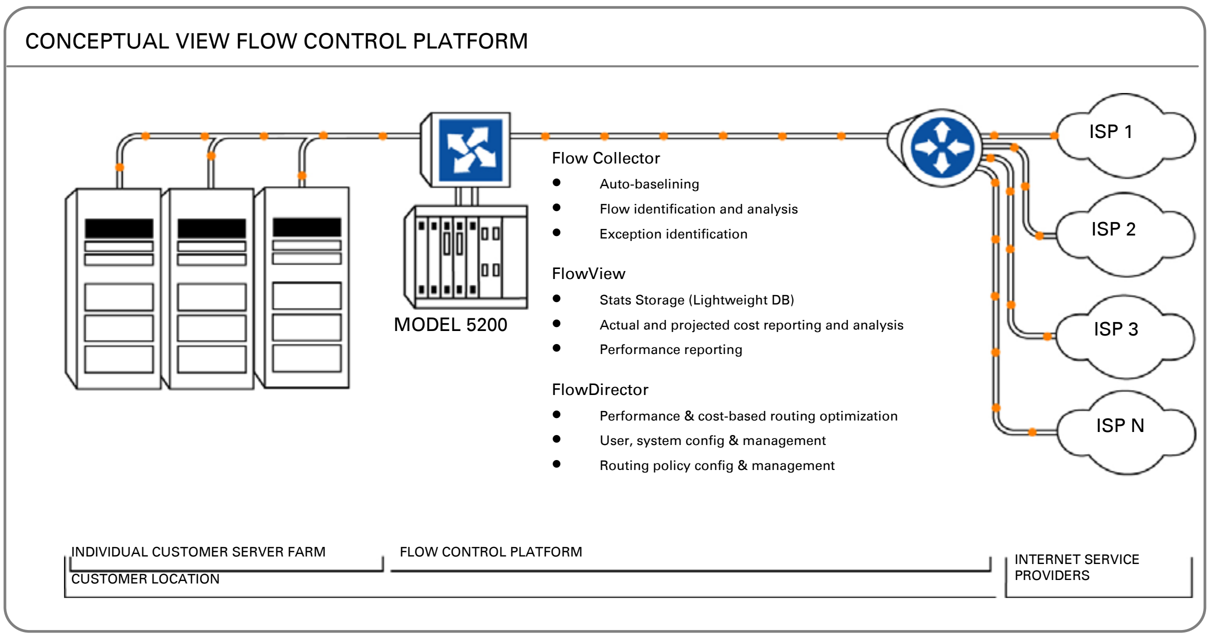

Figure 1: Flow Control Platform Architecture

The architecture also contains a usage and cost module based on the IntelliChoice model. This module is configured using customer billing contracts to produce provider cost thresholds. Performance thresholds can also be configured manually (including applications based), or they can be automatically determined (a unique process called automatic baselining) by the FCP solution.

With the information gathered by the Passive Flow Analyzer tool and the IntelliChoice technology, the FlowDirector™ processor is able to determine when traffic loads on any given provider exceed cost thresholds or when the performance to any given destination is less than optimal.

Single Variable Route Optimization

For many organizations, simply optimizing for one key objective across all traffic is a challenge that has not yet been overcome. Due to highly differentiated requirements, performance and cost optimization utilize different processes to meet its goals.

Performance-based route changes are the simplest and are made on a per destination basis. Usage levels on candidate providers and flow volume should still be considered, but in the end, performance-based decisions are made independently per destination and are not provider-based. Performance decisions are made in real-time, and do not require a database of information to determine the best path.

Such day-to-day performance gains can be substantial and ensure that Internet traffic performance remains consistent and predictable. But performance improvements are most significant when large-scale network outages occur. Under these circumstances, premise-based optimization provides business continuity by proactively routing around incapacitated providers or downstream network hotspots, delivering the best possible network performance for all destinations.

OPTIMIZING FOR PERFORMANCE

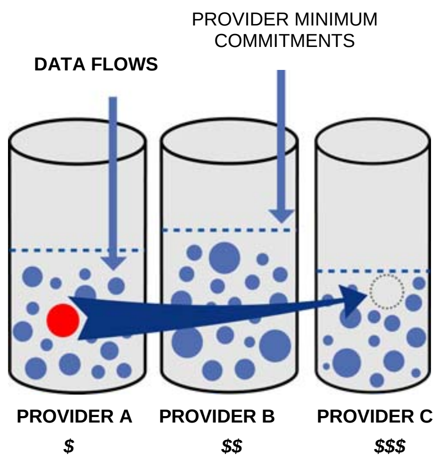

Figure 2: Performance Route Change

The Flow Control Platform solution's third-generation advances in performance optimization include scalable performance evaluation beyond Cisco NetFlow and active probing for key networks, as well as aggregated route change announcements that can proactively avoid performance issues. Although in-depth treatment of these topics is outside the scope of this paper, please refer to the white paper on *Premise-Based Optimization: How it Works* for more information.

In contrast, cost-based route optimization delivers the minimum network cost over time. Cost optimization is not a problem that can be handled on a 'per destination' basis, nor can minimum cost be enforced through accumulated decisions made for individual destinations. Instead, cost-based route optimization is an operation performed on a per provider basis. By understanding the current usage on a provider relative to the usage levels that guarantee the minimum cost, cost violations can be avoided.

When considering cost optimization, two concepts are essential to understand: minimum commitment and marginal cost of a data flow. Bandwidth contracts are generally structured with some minimum commitment and one or more burstable tiers, meaning that there is some base cost which the customer will always have to pay, and then an additional cost based on any additional usage. A minimum commitment is the bandwidth that an enterprise must pay for (and has at its disposal), regardless of how much it uses. Marginal cost of a data flow is the additional cost required to carry some added traffic load above the guaranteed minimum. Marginal cost is continuously changing. For example, if an enterprise has not yet met a minimum commitment, marginal cost of additional traffic that does not exceed the minimum commitment is zero. Traffic that achieves a new peak utilization beyond the minimum commitment of a provider contract has some marginal cost. Thus, marginal cost can be minimized by intelligent allocation of traffic flows in relation to billing profiles and peak utilization of each provider. This will be explained in detail in the next section.

Solutions that rank or prefer providers based on cost and then use that ranking to influence individual route-control decisions cannot ensure the minimum network cost. While individual route changes certainly can be made based on this ranking, there is no guarantee that the complete system of multiple providers with differing billing profiles will maintain minimum cost. Because most of these types of solutions are blind to the marginal cost of any route change, these solutions cannot guarantee that a given route change will not maintain the lowest possible marginal cost. Though such solutions may reduce some costs by maintaining preference for lower cost providers, the failure to enforce lowest marginal cost for the system prevents the strict minimization of bandwidth cost optimization.

IntelliChoice Minimum Network Cost

When a multi-homed enterprise runs BGP, the resulting traffic distribution across the service providers is determined using BGP's standard metrics. These metrics do not consider the cost or the current usage of the service providers.

As a result, the traffic is often distributed in a way that needlessly increases network costs since the bandwidth available is not efficiently used by BGP. In fact, because of BGP's preference for least number of router hops, it often selects the most expensive route to a destination. While Tier 1 providers often have the best peering, they are often the most expensive. Based on minimum commitments, burstable tiers and BGP's preference for the shortest route to a destination, we can create a maximum cost curve that mimics BGP's performance. To make matters worse, Tier 1 providers do not necessarily provide any guarantee that those peering points will provide a certain level of performance.

In contrast, the Internap FCP solution with the IntelliChoice technology leverages the diversity of multiple billing contracts to ensure the minimum dollar cost for any amount of usage. As described in the previous section, typical provider billing profiles involve a minimum cost commitment followed by a burstable tier at some dollar rate per megabit per second transfer rate. By understanding these contracts, and distributing load to achieve minimum marginal costs, it is possible to optimize the costs of overall bandwidth usage.

The table below outlines hypothetical contracts for three service providers, and includes current usage by a hypothetical enterprise customer.

Provider	Minimum Commit	Peak Burstable	Burstable Cost	Capacity Used
A	5 Mbps	50 Mbps	$50/Mbps	3.0 Mbps
B	10 Mbps	50 Mbps	$100/Mbps	7.0 Mbps
C	15 Mbps	50 Mbps	$150/Mbps	12.0 Mbps

Table 1: A Hypothetical Enterprise Customer

In the current situation, the customer has not yet met minimum commitments on any of its three providers. As such, any additional traffic added to any provider has a marginal cost of zero, so long as the commitment threshold is not exceeded.

While link preference models of route control may simply note that Provider A is the cheapest, the IntelliChoice model recognizes the sunk costs of minimum commitments. As such, it will seek to let the minimum commitments fill on all three providers before utilizing any burstable bandwidth. If a commit level is reached, traffic on that provider will be moved to another provider with available commit and equal performance. This process is sometimes referred to as commitment stuffing.

Next, let us assume we have met all minimum commitments. Examine table 2 below.

Provider	Minimum Commit	Peak Burstable	Burstable Cost	Capacity Used
A	5 Mbps	50 Mbps	$50/Mbps	5.0 Mbps
B	10 Mbps	50 Mbps	$100/Mbps	10.0 Mbps
C	15 Mbps	50 Mbps	$150/Mbps	15.0 Mbps

Table 2: A Hypothetical Enterprise Customer

In this situation, any additional traffic will force us to incur additional cost. Logically, the cheapest choice is to fill Provider A as much as possible, Provider B second and Provider C third. This produces lower marginal cost per megabit transferred.

Based on the combination of commitment stuffing and marginal cost analysis, we can create a curve representing the lowest possible cost of bandwidth at any given traffic rate.

Figure 3 on the following page illustrates the IntelliChoice model for minimum network cost in a multi-homed enterprise that has three service providers. This model calculates usage thresholds for each provider that must be enforced if cost is to be minimized. By continuously calculating and enforcing the usage levels through the IntelliChoice model, the FCP solution ensures minimum network cost.

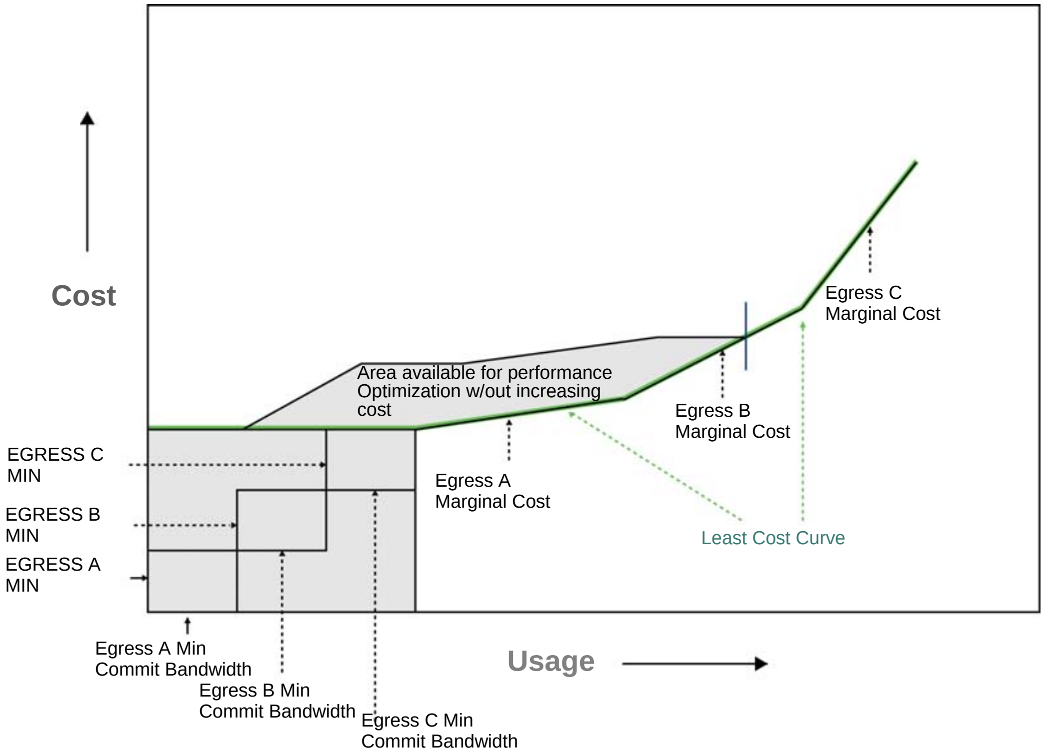

Figure 3: The IntelliChoice Model

The IntelliChoice model first fills all minimum commitments, as the additional cost to do so is zero. As usage increases, the IntelliChoice model fills each of the burstable tiers starting with the least expensive (A) and ending with the most expensive (C). Each burstable tier in the model is represented by a sloped line, with the slope relative to the cost of burstable bandwidth. By following this Least Cost Curve, minimum cost can be determined for any peak network usage. Simply draw a vertical line at the expected monthly usage (represented by the blue line in Figure 3) and the point where that line intersects the least cost curve is the provider usage threshold. The corresponding point on the y-axis represents the minimum dollar cost for that amount of bandwidth.

In contrast, we can overlay the Least Cost Curve of the IntelliChoice model with the Maximum Cost Curve produced by BGP as shown in Figure 4 on the following page.

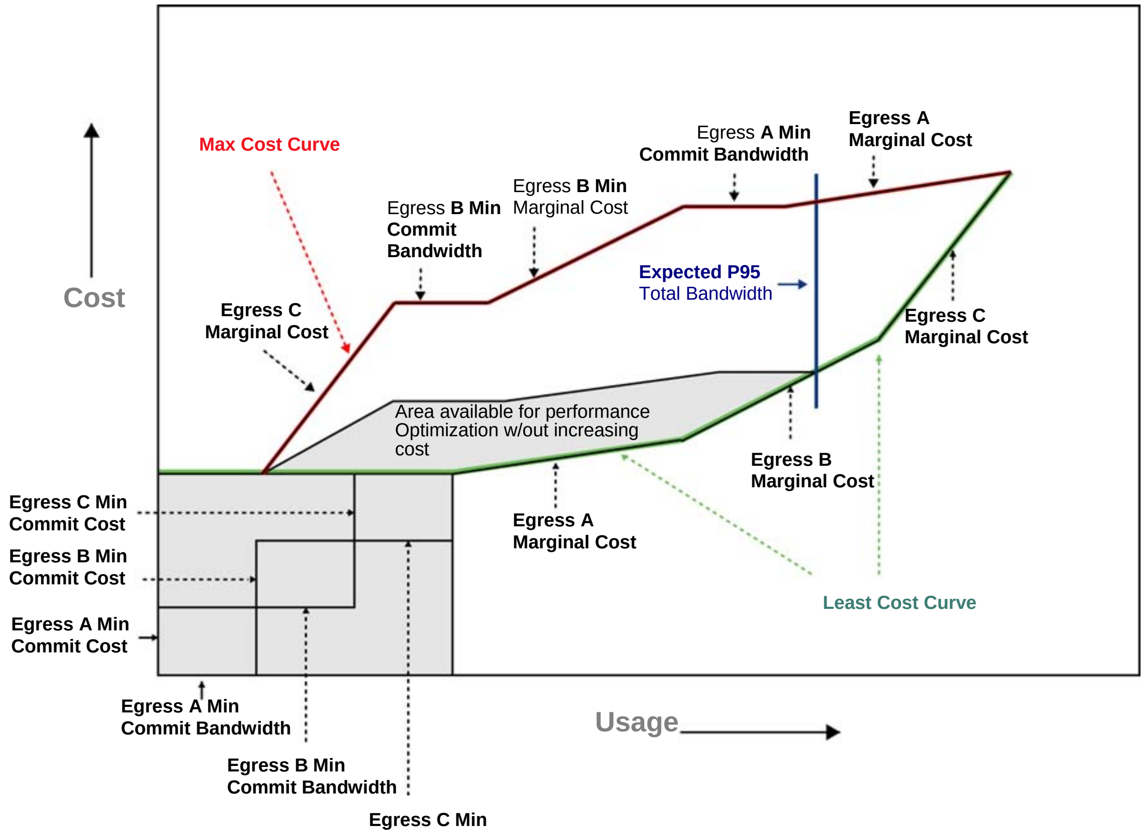

Figure 4: Least Cost Curve and Maximum Cost Curve

Once the IntelliChoice model determines cost thresholds, the FCP solution can monitor bandwidth usage levels for each provider relative to the defined thresholds. When usage approaches or meets a threshold, a cost violation is indicated and the FCP solution takes action by moving any additional traffic to next lowest cost provider.

The above description of the IntelliChoice model is relatively simple. There are two use cases, however, that present significant challenges to maintaining the absolute minimum bandwidth cost in a multi-homed environment. These two cases are described in the following sub-section.

Cost Grooming

As often happens in Internet environments, bandwidth utilization cyclically spikes and then enterprises are charged not on average usage, but peak usage (usually the 95th percentile of peak usage). Thus, once an enterprise hits a new peak usage on a particular provider, it is in the customer’s best interest to continuously use bandwidth at a rate as close as possible to that new usage rate in order to create the lowest possible marginal cost for that additional bandwidth. Let us create a new example in Table 3.

Provider	Minimum Commit	Peak Burstable	Burstable Cost	Capacity Used	Capacity Used (95th%)
A	5 Mbps	50 Mbps	$50/Mbps	5.0 Mbps	5.0 Mbps
B	10 Mbps	50 Mbps	$100/Mbps	10.0 Mbps	10.0 Mbps
C	15 Mbps	50 Mbps	$150/Mbps	15.0 Mbps	25.0 Mbps

Table 3: A Hypothetical Enterprise Customer

In this scenario, Provider C peaked at 25 Mbps at some time in the past, and has since subsided to 15.0 Mbps. While such a situation might not be the norm given the high cost of C, it could happen because of downtime on Provider A or B, or if performance requirements were such that Provider C was the only viable option for certain destination networks. A more detailed discussion of simultaneous cost and performance optimization is in the next section, IntelliChoice Simultaneous Optimization of Multiple Variables.

Under these circumstances, second-generation optimization devices would simply place additional traffic on Provider A. Any device that depends solely on the billing profile would make this mistake. The IntelliChoice technology, however, understands the state of each provider, and is able to use this information intelligently to minimize marginal cost.

Using the IntelliChoice model, the Flow Control Platform solution would allocate the next 10 Mpbs of traffic to Provider C. This is because, regardless of whether the enterprise uses Provider C at 25 Mbps or anything less, the enterprise will be charged at a rate of 25 Mbps. Anything above 25 Mbps in this scenario would then be placed on Provider A.

Another important corner case arises when an enterprise is close to cost thresholds on multiple providers as shown in Table 4.

Provider	Minimum Commit	Peak Burstable	Burstable Cost	Capacity Used
A	5 Mbps	50 Mbps	$50/Mbps	4.5 Mbps
B	10 Mbps	50 Mbps	$100/Mbps	9.8 Mbps
C	15 Mbps	50 Mbps	$150/Mbps	14.5 Mbps

Table 4: A Hypothetical Enterprise Customer

Although in this case all the thresholds in discussion are the minimum commitment thresholds, this is not a requirement for the use case. It could be similar to the previous example, where the new threshold was in a burstable tier.

If an additional 1 Mbps of traffic is requested from a new prefix address space (i.e., the aggregated requests of multiple users topologically close to each other), this traffic would exceed thresholds on any one of the available providers. In such a case, the Flow Control Platform solution performs a route swap. Certain flows equaling 0.5 Mbps in volume are selected from Provider A and moved to Provider C.

The newly requested traffic now flows on Provider A, and all cost thresholds are still met. In the case that cost thresholds cannot be maintained (e.g., 2.0 Mbps of traffic was requested in the above scenario), the FCP solution will perform route swaps to minimize the new peak utilization. This process, called cost grooming, maintains the minimum marginal cost, and hence, the minimum overall cost. Figure 5 illustrates the cost grooming process.

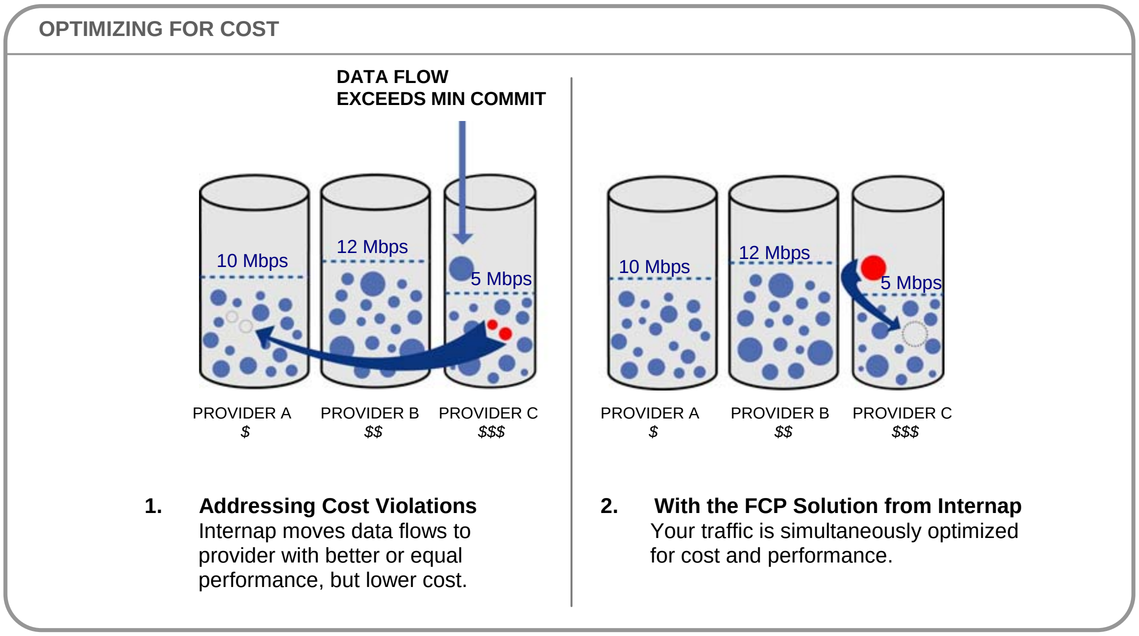

Figure 5: Optimizing for Cost

Such an action is only possible if the optimization system is fully aware of all concurrent traffic flows. The Flow Control Platform solution, utilizing the Passive Flow Analyzer tool, can handle over 100,000 destination networks to ensure optimized, minimum cost for an enterprise.

IntelliChoice Simultaneous Optimization of Multiple Variables

While the IntelliChoice cost grooming provides a significant, sustainable advantage over second-generation optimizatin solutions, most enterprises would like to maintain a certain level of performance while optimizing bandwidth costs. The algorithms to optimize performance and cost simultaneously for all destinations represent a significant challenge and require a tremendous amount of information regarding the number, size, performance and current route of every flow. In addition, such a solution requires the ability to consider the customer's provider contract structures. Utilizing the Passive Flow Analyzer tool for scalable information collection and the IntelliChoice model for advanced decision-making, the FCP solution is the only premise-based optimization solution on the market today that can simultaneously optimize for cost and performance.

This solution enables the administrator to set a primary and a secondary variable for optimization. There are two common cases through which the FCP solution enables simultaneous cost and performance utilization: (1) where the primary variable is equal across multiple providers, and the secondary variable is used to decide the route; and (2) when optimizing for one variable would cause a violation in the other variable.

In the simple case where the primary variable is equivalent on multiple providers, the secondary variable can be used to provide maximum user quality with no additional cost to the enterprise as shown in Table 5.

Provider	Minimum Commit	Current Capacity Used	Latency
A	5 Mbps	4.5 Mbps	100 ms
B	10 Mbps	9.0 Mbps	200 ms
C	15 Mbps	15.0 Mbps	60 ms

Table 5: A Hypothetical Enterprise Customer

Assume in the above example that cost is the primary variable and performance is the secondary variable (known as a "cost sensitive" policy). If an additional .5 Mbps were requested, it could be sent via Provider A or B with no additional marginal cost (as they are both under their minimum commitment). The FCP solution utilizes the secondary variable, performance, to select Provider A. As described in the section above, continuous cost grooming occurs to maintain the absolute lowest cost. If there were a provider who clearly offered a lower marginal cost for additional traffic, that provider would have been chosen. Decision flow for a performance sensitive policy is similar under normal circumstances.

Nevertheless, performance-constrained situations still occur, where one or more of the candidate providers are at or near their cost-based thresholds. To resolve these performance problems normally would induce a cost violation, and similarly to mitigate the cost violation would not necessarily repair the performance problems. However, even in these situations, the FCP solution can utilize the superior visibility that the Passive Flow Analyzer tool provides to perform a cost/performance swap. This swap is performed similarly to the swap described in the cost grooming section above. To resolve the performance problem in a cost-constrained environment, the poorly performing prefix is swapped out with an equally sized prefix (or collection of prefixes) on another provider. The prefixes to be swapped are selected because the performance between the two providers is equivalent or better for the given destination networks, thus preventing a subsequent performance problem. This premise-based optimization method allows the FCP solution to maintain guaranteed minimum performance levels while optimizing costs.

In the uncommon case when the FCP solution is so constrained, either because there are only a small number of flows, or the providers are aggressively diverse, a policy setting is available that enables the FlowDirector processor to select between a cost- or performance-based move for a given destination. This acts as a final arbiter for a situation that rarely occurs.

Figure 6 illustrates such a situation where the FCP solution is cost constrained and a performance problem occurs. The example illustrates how the FCP solution is able to simultaneously resolve the performance problem without increasing costs. Because of the visibility that the FlowCollector processor provides into every flow, the FCP solution finds a candidate prefix that can be swapped and resolves the problem without incurring additional cost. The performance problem on Provider A can only be solved by moving the flow on one of the two cost-constrained providers, B or C. This would increase the marginal cost of Provider B or C beyond the minimum (as shown by the thresholds). However, the FCP solution has visibility into the performance of the additional flows of Providers B and C and is able to determine a prefix that would experience equivalent performance on Provider A. These two prefixes are then swapped by the FlowDirector processor in a move that preserves the minimum network cost and resolves the performance violation on Provider A.

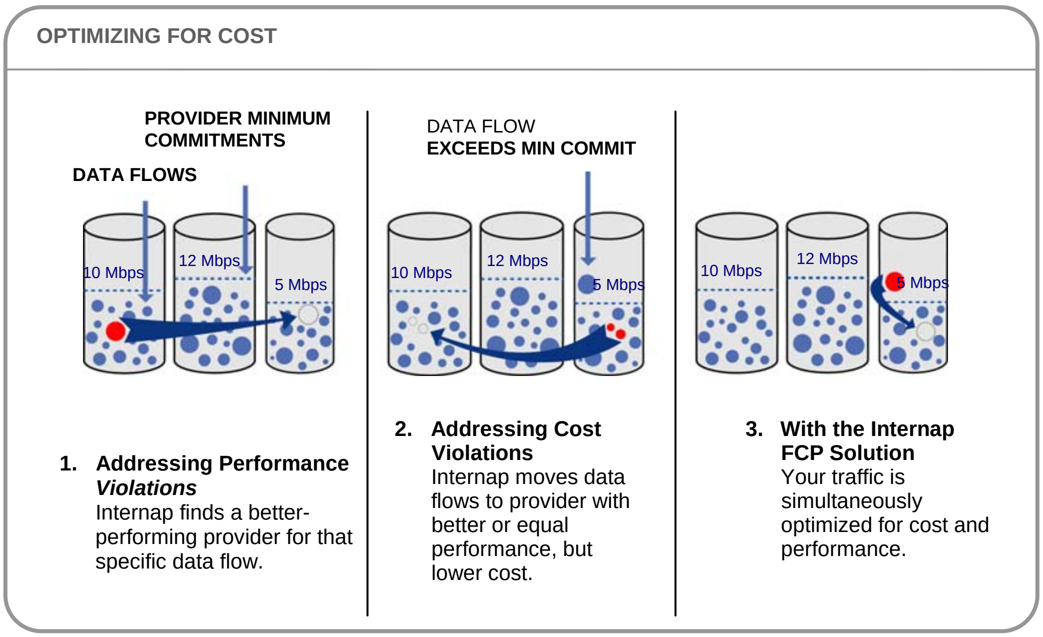

Figure 6: Optimizing for Cost and Performance Simultaneously

Summary

While previous generation premise-based optimization devices could provide marginal performance gains or cost reductions, an extensive set of variables must be monitored and managed to provide optimum cost reduction and performance enhancements simultaneously.

The Internap Flow Control Platform solution, with the IntelliChoice technology, can build a true least-cost model. Employing massive scale for monitoring all traffic, the FCP solution can then adhere to that model using route changes and route swaps between providers.

Combining the IntelliChoice technology with the Passive Flow Analyzer tool, the FCP solution can uniquely provide simultaneous performance and cost optimization. By understanding all current traffic flows in relation to provider usage and billing, the FCP solution can make cost/performance swaps that improve the performance of individual destination networks while optimizing cost determined by the IntelliChoice technology.

The Internap FCP solution is the only premise-based optimization solution that can simultaneously meet cost and performance goals for enterprises. Utilizing the FCP solution, enterprises can enable a reliable, consistent Internet experience for their customers and employees, while optimizing network costs.