

pulse networks
a real-time path to profitability



Introducing. . .

The Real-Time Interconnect Server:

**Dittberner's Assessment of a New Solution for
On-Demand, Cost-Adaptive Traffic Routing**

The Telecom Cost Management Race

The term "slim margin of victory" took on a whole new meaning in the 50th running of the Daytona 500 in 2008.

Ryan Newman, the winning driver in the most prestigious stock car race in America, took the checkered flag only 9/100ths of a second ahead of the second place car.

Nine one hundredths of a second is an awfully slim margin for a 3 hour and 17 minute race.

So what spelled the difference? Was it Newman's driving skills? Was the turbocharger in his Dodge better designed? Or was the race won in the pits because Newman's pit crew was slightly faster at refueling and changing tires?

Actually, the technology and process fine-tuning required to win the Daytona 500 are similar to what we go through in telecom cost management.

Are my operations streamlined to maximize revenue? Am I earning the very best margins from my interconnect partners? Are the processes I put in place 5 years ago still optimal for today's climate? -- Or are they slowing me down like an inexperienced pit crew?

These are key questions telecom executives must ask themselves over and over again, especially during tight economic times.

The Purpose of This Paper

Hello, and permit us to introduce ourselves. We're the analyst team at Dittberner Associates, a market research firm that's been following the telecom revenue management scene for several years now. In that time we've published several research studies on topics such as telecom billing and charging, revenue assurance, fraud, and cost management.

You're reading this paper because Pulse Voice Inc, learning of our specialty, hired us to interview their customers and write this paper to spread the word about a new solution it's developed to help telecoms better manage their interconnect costs.

The Telecom Revenue Management Market

It's no secret that in the last several years, the risks of operating a telecom company have increased substantially. Cable, internet, consumer electronics,

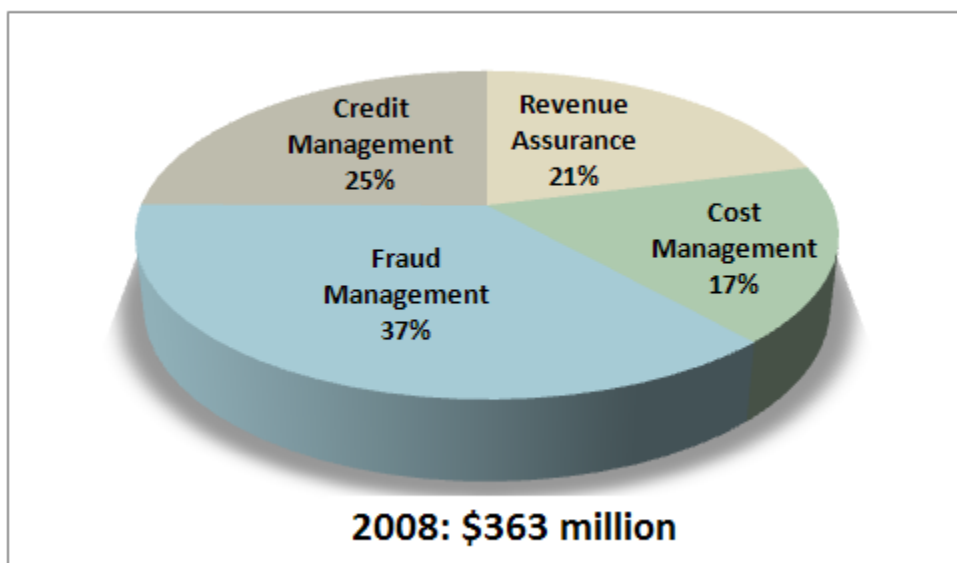
and computing companies are competing for a share of communications revenue pie.

And in an environment like that, picking winning services has become only a little bit more certain than rolling the dice. Because it's so hard to accurately predict service winners, telecoms are forced to become experts in rapid experimentation, and that implies having enormous operational dexterity. Telecoms must not only launch services quickly, they must also bill for them creatively and provision them efficiently to keep ahead.

Which brings us to an interesting conclusion. At a time when operational systems and OSS/BSS platforms are being stretched to their limits, it's crucial that revenue leaks, fraud, customer bad debt, and interconnect billing problems don't sink the business in the midst of all the system juggling.

It's a key reason there was a robust, \$363 million market for telecom revenue management solutions in 2008, comprising four key areas: fraud management, revenue assurance, cost management, and credit management. See chart below.

The Telecom Revenue Management Solutions Market



Source: Dittberner Associates

Why Telcos Have Trouble Managing Their Interconnect Costs

Dittberner categorizes “least cost routing” as a cost management activity, putting it in the same family as functions such as interconnect bill verification and content management assurance.

Of course, least cost routing is a rather unique cost management function because its job is to assure revenues ahead of time, rather than determine what went wrong in the billing process. In this sense, least cost routing is a kind of “pre-cost management” activity.

However you choose to classify least cost routing, you’ve got to admit that it’s an assurance activity that’s still maturing. Practically all carriers have “least cost” or “best value routing” programs in place, but unfortunately most carriers are barely keeping up, especially now that some interconnect suppliers are updating their rate decks on a weekly, sometimes daily basis.

For years the suppliers of bandwidth, long distance and roaming services have been in the driver’s seat. The key problem is it’s cumbersome to optimize routing lists for all suppliers to all destinations. One supplier may have fabulous rates to Milan, Italy or Grand Rapids, Michigan, but if its rates are only cheap to one destination, then you can’t affordably use them because it means redoing your entire routing tables for only one termination point.

In short, if you’re racing to calculate your routing policy, you need to avoid costly pit stops.

The Traditional Method of Best Route Calculating and Switch Translation

OK, so let’s briefly look at why the traditional approach to least cost routing is so cumbersome.

After interconnect suppliers send you their rate decks, it’s usually your business unit’s responsibility to prioritize traffic routes and implement policies based on lowest cost, highest quality, and other parameters. However, the rate decks arrive on irregular schedules and in a variety of Excel/database formats, so there are delays in merging the data and creating an optimal routing plan.

Once the optimized routing instructions are created, they’re sent to engineering where the job gets even tougher, for it’s engineering’s job to translate that policy into specific routing instructions for a diverse network of proprietary switches whose routing capabilities are rather weak.

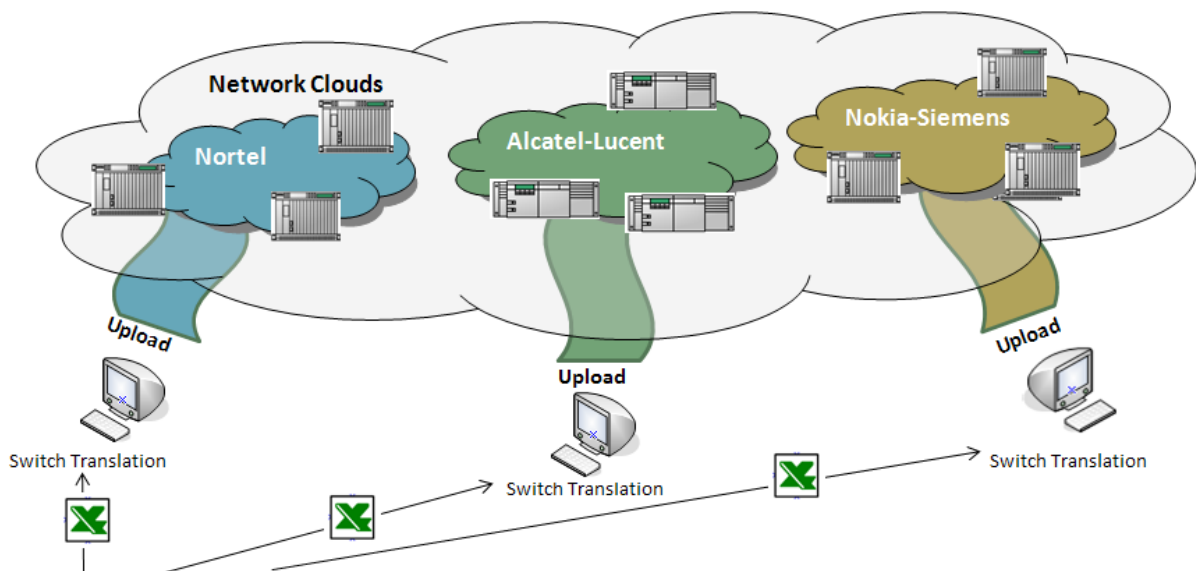
The chief switching issue is the lack of online database capacity. For example, engineering may need drastically to condense the 2 to 3 million records in the routing policy to the capacity a switch can handle. Some Nortel switches can

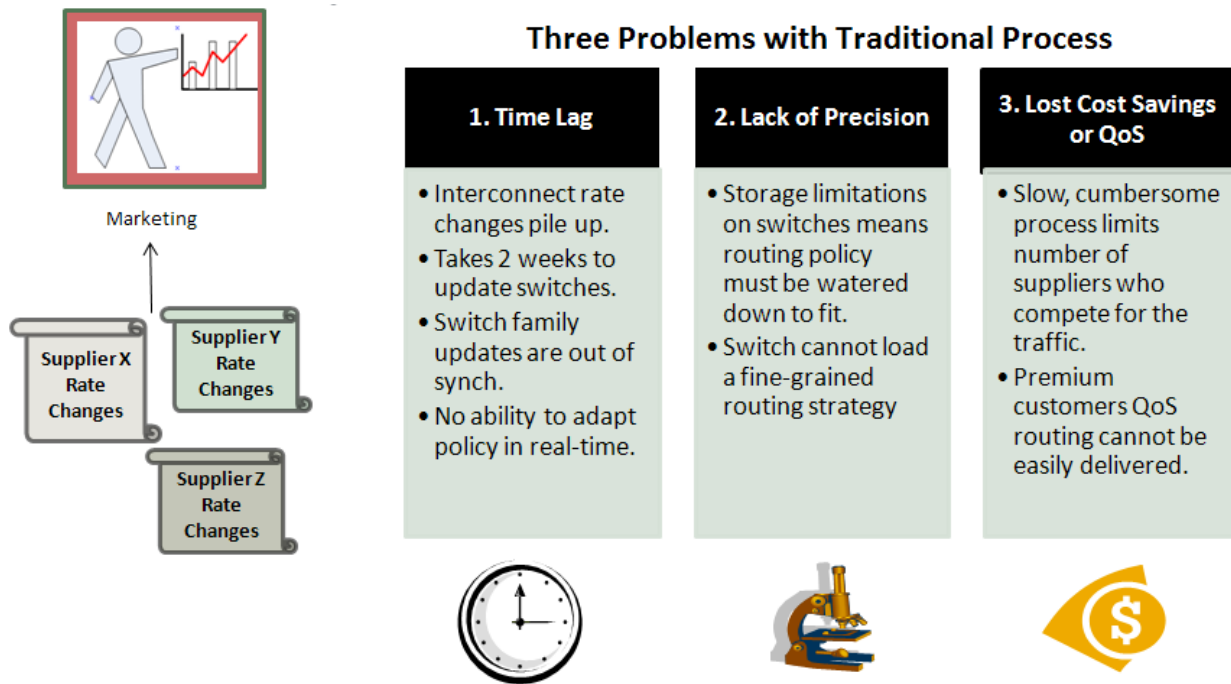
accept only 200,000 records. Even the latest generation softswitches often accept a maximum of only 500,000 to 600,000 records.

Now, since your engineers are forced to squeeze routing data down to fit each switch's storage limit, they have no choice but to sacrifice routing granularity, optimization, -- and ultimately cost savings.

Generally, the job of putting routing policies into the network requires the work of several switch translation engineers. Depending on the switch diversity of your network, we estimate a full-time translation engineer is needed for every 5 to 10 switches. Of course, as you move to more next-generation and IMS-flavored networks, calculating your routing policy gets even more complicated, demanding even more experts on staff. The figure below summarizes the problems with the traditional methods of least cost routing.

Traditional Least Cost Routing Process





The Need for an Automated, Real-Time Solution

Given what we've just discussed so far, is it any wonder that current least cost routing processes can't keep up?

The irony here is that, because the process is so unwieldy, you -- the customer who's buying interconnect traffic -- are not really in control. And the only way to gain control is to find a more real-time method of updating your switches.

When you can achieve that, instead of terminating your traffic with only a few suppliers, you can maintain agreements with a dozen or even two dozen interconnect partners, and thereby introduce some old-fashioned competition. Suppliers will be bidding against each other for your traffic, and you'll be able to cherry-pick the best ones you want for the destinations you need.

At that point, your business units will finally be able to drive volume deals and bring advanced purchasing techniques into play such as routing for peak rates and controlling your capacity and costs like an airliner controls the prices of the seats on its jets.

Introducing the Real-Time Interconnect Server

The rosy picture we've just painted is no longer wishful thinking. Some innovative carriers have successfully tackled the switch translation problem and transformed their least cost routing solutions into major money-savers.

As strange as it may sound, the answer to the switch translation problem is to largely get out of the switch translation business.

The new way to do that is to employ a real-time interconnect server attached to the network whose job is to route traffic for all switches via multiple types of signaling. Every call that requires routing – local, long distance, and international -- dips into that server, which finds the optimal route, and returns instructions to the switch on where to send traffic.

The beauty of centralizing routing policy on a standard server database like this is that you can make your routing data and policy as granular as you'd like. At the same time, the need for switch translation work largely goes away.

Frankly, it's puzzling that a solution like this has taken so long to come to market because attaching servers to networks is as old as 800-numbers, IVRs, and Caller ID. Switches were never designed for sophisticated call routing, so it makes perfect sense to offload that work to an intelligent network platform.

Real-Time Interconnect Pioneer - Tier 2 Wireless Carrier

One of the first carriers in North America to employ a real-time interconnect server is a Tier 2 Wireless Carrier, where the solution has been working successfully now for two years.

This Tier 2 Wireless Carrier is an operator in the U.S. with about 2,000 employees and 3.6 million subscribers/pre-paid customers.

The Carrier's cost management challenge was as simple as it was hard to deliver: pare down interconnect costs so it could stay profitable in a brutally price-competitive U.S. wireless market.

Making cost cutting even more critical was the company's policy of offering customers a bucket of unlimited local calling and text messaging for a low monthly fee. Customers could also set up pre-paid long distance account that's debited for better expense control. Finally, the Carrier offered an unlimited local and long distance plan for calls within the U.S. for a fixed fee.

The real-time interconnect server has been instrumental to this Carrier's program of lowering costs and managing a greater number of interconnect partners. The solution has kept up with the load despite the fact that the pace of required updates has increased substantially. Three years ago, most rates changed semi-annually, but today those rates often change weekly.

Every month, the Carrier's product managers create a new set of rates and routing policies. To implement that policy is now quite simple: they go to a website, upload the rates and the interface to the switches is done automatically.

In the past, engineers required about 2 weeks to put those rates into a queue. The data was then loaded, usually at night, to the Tier 2 Wireless Carrier's 30 switches so that in the event anything went wrong, it could revert back to the prior version. Another difficulty was that not all switches were loaded at the same time because different families of switches were on different time tables for updates.

Today, as calls come into the Carrier's switches; the switches are programmed to say: "Don't use your internal tables. Instead, go the master switch (which is fed by the real-time interconnect server) to get your routing instructions."

As this Carrier's experience with the platform has grown, it's steadily made its policy more granular. Today, routing is based on several parameters, including: cost, quality of the call, time of day, and long distance jurisdiction such as LATA/OCN. They can also trigger changes based on carrier commitments to send a certain amount of traffic at a certain rate.

Real-Time Analytics for PRIMUS Canada

PRIMUS Canada is a multi-service operator serving 1 million customers and offering a wide selection of consumer and business services, including local, long distance, TalkBroadband (VoIP), wireless service, and high speed DSL.

Two years ago, PRIMUS replaced its older generation Nortel DMS-based switches with next-generation switches from Genband. Along with that change, it decided to migrate from distributed to central routing using a real-time interconnect server. Today, eight switching gateways are served by the real-time platform.

For PRIMUS, a key benefit of the real-time interconnect server approach was its ability to measure the gap between the "as written" versus "as delivered" routing policy.

Switch and trunk overloads are a common occurrence, and so PRIMUS discovered that it doesn't make sense to send all your calls to the least cost route if 70% of those calls don't get through or the caller hears only a fast-busy signal. That's a recipe for frustrating customers.

Here's where the real-time CDR collection of the real-time interconnect server proved vital: in a matter of minutes, the server could calculate the percent of time a call sent via a low cost route actually terminated on that least cost route.

That information wasn't visible before. Yes, you could block calls to an overloaded trunk, but you couldn't reroute to an optimal second, third, or Nth routing choice. As a PRIMUS executive put it: "In those days, we were merely capacity managing our network, not cost managing it."

So how did PRIMUS detect low call delivery ratios in the past? It either waited for yesterday's reports or heard about it from the call center because customers were complaining.

Now, leveraging its real-time server, PRIMUS has fine-tuned its least cost routing policy to route traffic to or away from suppliers based on the desired trade off of least cost v. call completion for particular customer groups.

Advantages of a Real-Time Interconnect Server

OK, we've reviewed how two telcos are applying a real-time interconnect server for competitive advantage. Let's now discuss eight major benefits of the approach:

1. Savings from Optimized Routing

While it's tough to precisely predict what savings you'll earn by deploying a real-time interconnect server, through an ability to support a greater number of suppliers and the greater granularity of traffic routing the server delivers, we estimate interconnect margin savings as high as 9%.

2. Deep Routing Granularity

With a server in control of your switch routing, your days of trying to force-fit policy into low capacity switch databases are over. No longer do you need to sacrifice the granularity of your policy to meet switch limitations.

Here's an example of where this greater degree of routing granularity can pay off.

Say you need the least cost route to a particular neighborhood of San Diego. If your routing policy isn't fine grained enough to terminate on the lowest cost supplier in that neighborhood exchange, then you're required to pay the higher **average rate** offered by all carriers serving that neighborhood.

The real-time interconnect server gives you the power to achieve full 10-digit routing because you're no longer limited by the routing storage capacity of your switches.

3. Adaptive Cost Routing

Speed is an important aspect of your routing policy. If an interconnect supplier offers you a splendid volume discount, you'd like to route traffic to them right away. However, if it takes your translation team two weeks to upload interconnect rates into the network; you've lost an opportunity to save.

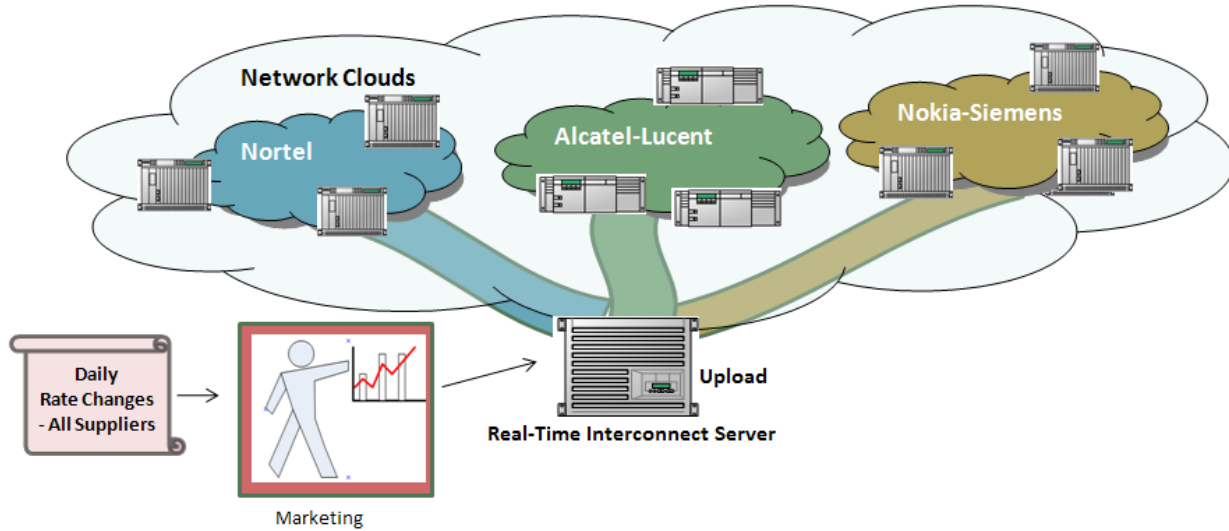
With a real-time server, you're no longer forced to route traffic based on yesterday's routing table. If calls are not getting through to your first choice destination, then the system dynamically provides alternatives based on real-world conditions.

In this way your interconnect server mimics the adaptive routing of network traffic to avoid congestion points, except that in this case, you're doing **cost adaptive routing** – making changes either to lower costs or increase the business value of your interconnect flow.

Minutes after your calling patterns start to change, the server is changing its decisions based on the criteria you set for your expectation of call quality by a particular customer group.

This is a big improvement over the world of static routing tables sitting on switches. Your central server is constantly analyzing events and saying: "Hey, my situation changed 10 minutes ago, so I'm now altering the traffic flow . . . to save money . . . achieve greater call completion . . . or to optimize policies X, Y, and Z."

Least Cost Routing with a Real-Time Interconnect Server



Three Key Advantages

1. Suppliers Compete for Your Traffic

- Time-consuming switch translation process is greatly streamlined.
- Marketing writes its policies directly on the server.
- Backlog of rate changes goes away, enabling carrier to work with more suppliers.

2. Implement Fine-Grained Routing Policy

- Switch storage limitation goes away, enabling a rich and flexible policy.
- Policies are centralized and not "lost in translation".
- Choices multiply in terms of suppliers, rates, time-of-day routing, and QoS.

3. Fast & Adaptive Cost Management

- Changes are made on-demand to save money, improve QoS, or other policy.
- Greater speed allows greater procurement sophistication and route optimization.
- System is programmed to automatically adapt to real world conditions and opportunities.



4. Real-Time Interconnect Fraud Control

Being able to centralize data on interconnect calling allows you to catch fraud in minutes as opposed to days. Perhaps a calling card has been compromised. Or maybe there's been a PBX security breach so someone is redirecting calls to Cuba or the Middle East where international calls are \$3.00 a minute.

Most telecoms have dedicated experts whose job is to read reports on the behavior of high volume callers and detect these things. Problem is, it might take them two or three days to detect a major issue.

However with an interconnect server collecting CDRs for you in near real-time, you're able to quickly detect the fraud, change the traffic flow, and redirect "compromised" accounts to customer service. A company may be billing \$1,000 of traffic in a month, but in a couple of hours, the server raises a red flag because it detects calling at the rate of \$10,000 a month.

The benefit here, of course, is that you're saving your customers from headaches and you no longer have to hear complaints like: "How come you couldn't stop this faster?"

5. Better Quality Service for Key Customers

Because you can set a detailed routing policy around quality metrics, it's far easier to deliver multiple tiers of service to your customers. The server, knowing the customer source of each voice call, can automatically direct calls to premium or lower quality routes as your business policy requires.

For customers subscribing to your economy brand, the routing parameters you specify are based on relatively poor quality, but low price. Conversely, you define a high quality routing policy for your premium corporate customers.

So rather than maintain separate routing tables and equipment for different sets of customers, your business strategy is implemented in one database that you can change on the fly.

6. Pre-Cost Management

Executives responsible for improving bottom line revenue know that a bird in hand is worth two in the bush.

So when you can drive down your interconnect costs by stimulating competition, you earn a favorable margin ahead of the game. In that way you're driving a two-pronged cost management strategy: scrutinizing interconnect bills at the same time you're proactively getting better margins from your suppliers.

In fact, the platform can effectively "train" your interconnect partners to treat you right. You're in control now because you can direct your traffic to whatever partner is giving you the best rates, quality of service, or even interconnect billing support.

7. Applying Local Number Portability Information

A real-time server can also apply local number portability data to routing. The database dip tells you if it's a ported number and if it is, then the new local routing number (LRN) is added and routed. In the U.S. market, you save money on the porting fees that interconnect suppliers charge.

8. Reduced Labor and an End to Tedious Translation Work

With a real-time server in place, the need for switch translation is considerably reduced. Whilst it's prudent to maintain a strong switch translation expertise for backups, emergencies, and special situations, most of the engineers can be assigned to higher priority work.

About Pulse Voice Inc.

pulse has been providing world class communications solutions globally since its inception in 1996. pulse networks, a division of Pulse Voice Inc., through its global offices in North America, Europe and Asia; offer a portfolio of real-time margin assurance and intelligent network solutions to carriers and services providers of all sizes. pulse networks portfolio is heavily driven by Return on Investment (ROI), payback and cost reduction strategies, ensuring that any investment made in a pulse solution has a significant and immediate impact on a carrier's operating costs. Currently pulse Networks has implementations in over 15 carriers and service providers worldwide that have benefited from these savings.

In addition to the Network division, pulse maintains an Enterprise group that specializes in automated interactive call center technologies to an international client base that includes Fortune 500 companies and governments at the national, state and local levels.

With over 200 customers in 40 countries worldwide, pulse has a significant base of client advocates for its products and solutions.

pulse's real-time interconnect server is delivering value for several service providers in North America and Europe. And the benefits of its solution are by no means limited to long distance and international voice traffic routing. The pulse server can add real-time control to networks that employ SS7, SIP, CDMA, IMS and GSM protocols.

The pulse realtime server is generally useful in any telecom architecture where there's value in maintaining a single, central resource of data available in real-time.

A Latin American wireless operator, for example, is using the server to integrate GSM and CDMA networks under a centralized billing platform. Other uses are found in fraud management where the server maintains a black list of numbers to cut off calls or flag calls for routing to a call center.

Mobile operators and MVNOs in Europe are using the pulse platform to manage call control, authentication, roaming, pre-paid vouchers and online account balances.

Yet another example is number throttling. If you're getting flooded with calls because of a local emergency or people watching TV are calling in to vote for their favorite American Idol, then you can throttle the down to a maximum number of calls to prevent your network from being saturated.

Executive Summary

Best value traffic routing is a crucial cost management program, yet most service providers can't fully optimize their policy or achieve the lowest margins possible because they're hamstrung by cumbersome processes and switch limitations. The chief bottlenecks are these:

- Difficulties around the creation and switch-translation of routing policies force carriers to work with only a handful of interconnect suppliers; limiting competition and keeping interconnect rates artificially high.
- Interconnect suppliers who used to keep their rates constant for months at a time are now changing their rates weekly, sometimes daily. This puts pressure on carriers to quickly revise their policies and translate that policy to take advantage of the best rates.
- Legacy and next-generation switches lack the online database capacity to support robust routing policy, meaning the best a service provider can deliver is a diluted, sub-optimal policy to their networks that's at best two weeks old.

To solve these routing policy problems, this paper discusses the benefits of deploying a real-time interconnect server whose job is to centrally route traffic for all switches via signaling. With centralized routing off a standard database, policy can become highly granular and optimized with changes made on-demand. In addition to lowering interconnect costs by expanding competitive bidding among suppliers, the solution adapts routing policy to real-world results,

can deliver a stratified QoS policy by customer value, manage fraud in real-time, and reduce switch translation labor costs.

A note from pulse networks

We thank you for taking the time to read this Dittberner analysis of real-time interconnects servers. If you'd like a printed version of the paper, please contact sales@pulsenetworks.com or at 905-754-4100.

If you are interesting in learning more about pulse's real-time interconnect server, please visit www.pulsenetworks.com or contact sales@pulsenetworks.com or at 905-754-4100. We would be delighted to meet with your team to discuss how a realtime interconnect solution can be customized to your company's needs. Typical payback/ROI periods for our solutions are in the range of 6 months to a year.