

LB3 and TC2's **2014 Year in Review**



Telecom stories that will affect enterprises in 2015



In the world of what we are learning to call ICT (Information and Communications Technology), 2014 was like the years just before it – meaning that the only constant was accelerating change.

AT&T and Verizon kept pushing the end of TDM and the migration of the public network to IP. Frame Relay and ATM are gone, and other platforms that users have relied upon for decades (like AT&T’s VTNS) are following suit. If you are interested, see the pieces within by LB3’s Deb Boehling and TC2’s Janis Stephens on AT&T’s new Withdrawal of Service Matrix. And while you’re focused on carrier games, don’t overlook the piece by LB3’s Laura McDonald and Deb Boehling on Verizon’s [not-so] Rapid Delivery platform and the risks it creates for large users.

SIP trunking continues to advance – we rarely see major enterprise voice procurements that don’t include it (in Europe as well as the U.S.) and “best practices” for deploying it are beginning to appear. These are discussed in three articles by Jack Deal and Ben Fox of TC2, and Hank Levine of LB3, that appear in the pages that follow.

With SIP trunking has come a growing demand for Ethernet access, which continues to be challenging because even though more and more commercial buildings are pierced by fiber, overall penetration is still only around 50%. It is therefore almost certain that at least a few key customer sites won’t have fiber access, triggering dreaded special construction and attendant costs and delays. If you’d like to know more, we’ve included an article by Jack Deal and Hank Levine on this.

If SIP is no longer cutting edge in ICT, what is? Among the strong candidates is the cloud, and as cloud deals get larger and more ubiquitous the need to develop better terms is becoming acute. To help with this you’ll find two pieces by TC2’s Ben Fox and LB3’s Marc Lindsey on negotiating private cloud transactions in the collection.

Wireless is hardly cutting edge any more, but its rapid growth continues to be a major story. We’ve included a piece by TC2’s David Rohde that explores how the soaring demand for, and price of, spectrum is entrenching Verizon and AT&T as market leaders and frustrating competition in the wireless market.

There were two big regulatory stories in 2014, one of which – Network Neutrality – got all of the headlines. The FCC’s second attempt to codify Net Neutrality was struck down by the DC Circuit in January of 2014, but the agency’s response – to the chagrin of some and the delight of others – was to extend Title II regulation to Internet access (though not, contrary to Congressional teeth gnashing, to the Internet itself). LB3’s Colleen Boothby dissects the new FCC decision, and what it means for enterprise users, in a piece reprinted below.

The regulatory story that didn’t get nearly as much attention was the FCC-as-pit-bull. The industry was pleased (and users were nervous) when Tom Wheeler – who had served as President and CEO of both the National Cable Television Association and the Cellular Telecommunications & Internet Association – was chosen by President Obama to Chair the FCC. But he wasted no time in assuming the mantle of new Sheriff in town, not only on issues like Net Neutrality but also by encouraging aggressive action by the Enforcement Bureau with moves such as establishment of the E-rate Strike Force and imposing record fines on major carriers for cramming and security breaches. Most of these are of greater concern to carriers than customers, but not all – as evidenced by Steve Rosen’s article on the Commissions campaign against eRate waste, fraud and abuse.

In the world of M&A, the story was one of blockbuster acquisitions foiled but smaller strategic acquisitions going through – such as Level 3’s purchase of tw telecom. We’ve included a piece by David Rohde on Level 3’s turnaround and its implications for the market.

As the leading law firm (LB3) and consultancy (TC2) representing Fortune 100 companies in their procurement and benchmarking of network and related IT services, before the FCC, and in disputes with their carriers, we have worked on, and written about, all of these issues. This booklet brings together some of the articles we wrote on these (and other) subjects last year, and offers some thoughts on their implications for 2015 and beyond. We hope you find the collection useful, and welcome your comments.

Sincerely,



SIP Trunking

Enterprise Sourcing of SIP Trunking Services
Hank Levine and Jack Deal.....1

Three Things You Really, REALLY Need to Know Before You Tackle SIP Trunking
Hank Levine.....3

SIP Trunking is Driving Users’ Ability to Consolidate Vendors in Europe
Ben Fox.....4

Regulatory Developments

What the FCC’s “Open Internet” Ruling REALLY Means for the Enterprise
Colleen Boothby.....5

How to Avoid eRate Rule Violations
Steve Rosen.....23

New Technology

Negotiating Private Cloud Transactions: Building the Foundations of Best in Class Deals Part I Introduction and Pricing
Ben Fox and Marc Lindsey.....7

Negotiating Private Cloud Transactions: Building the Foundations of Best in Class Deals Part II Performance, Control, and Risk Allocation
Ben Fox and Marc Lindsey.....10

Why Ethernet Access is a Critical Part of Enterprise Wireline Networks
Hank Levine and Jack Deal.....13

Carrier Games

AT&T Issues Withdrawal of Service Matrix – Part 1
Deb Boehling and Janis Stephens.....15

AT&T Issues Withdrawal of Service Matrix – Part 2
Deb Boehling and Janis Stephens.....16

Verizon Rapid Delivery - Not SO Fast
Laura McDonald and Deb Boehling.....18

Level 3’s Turnaround and the Search for the Next First-tier Provider
David Rohde.....20

Gigantic Spectrum Cost is Making Wireless Commoditization a Pipe(-less) Dream
David Rohde.....21



Enterprise Sourcing of SIP Trunking Services

Hank Levine and Jack Deal

Over the last few years, large enterprises have begun to reap the benefits of moving from “traditional” TDM technology—PBX trunks and the like—to SIP trunking for voice services. SIP, if anyone doesn’t know it by now, is basically a way to provide voice as an application over IP data networks. To be precise, traditional TDM trunks use dedicated bandwidth for each call, and the call control is provided by in-band signaling or, if the trunks are ISDN trunks, the “D” channel. SIP (Session Initiated Protocol) provides a standardized signaling protocol for Voice over IP calls.

LB3 and TC2 are doing a lot of transitions from TDM to SIP—it’s moved from “cutting edge” to “off the shelf” and is now part of pretty much every Fortune 100 network RFP. Like all technology transitions, it’s complicated. When considering a move from “traditional” voice to SIP trunking, enterprises may find the following of interest.

Technical Complexity

Part of the complexity is technical. Service Level Agreements (SLAs) for SIP are one obvious issue. Sizing access pipes (SIP typically rides on an existing MPLS network) is another. Then there’s how many trunks you need to handle busy hour traffic (remember Erlang-B tables?).

Pricing and Rate Inconsistencies

Another issue that gives users pause are the pricing/rate inconsistencies between SIP trunk providers, both in the trunk (i.e., the concurrent call path) and the related minutes of use elements. Like MPLS in the early days, there is no uniformity in approach, so a fair amount of inventory and bid analysis is required to sort out the real cost of different vendors/bids. For example, do you need failover capability from one data center to another in a centralized deployment (this is both a technical and pricing issue)? How about DID numbers for each remote location? The transfer of toll free calls from SIP termination to another SIP termination (or to a legacy TDM location)? Each supplier treats these elements differently.

Contract Issues

Then there are the contract issues. The biggest up-front problem is making sure that your existing contracts don’t restrict, or “tax,” the transition, either because of restrictions on migration or because of early termination or shortfall charges that you may incur if you move “too soon.” AT&T (which has resisted SIP far more than Verizon) has been particularly bad in this regard. Even as it loudly proclaims the need to phase out legacy PSTN services, AT&T often structures its legacy local services agreements to prevent users from moving away too quickly (think “managing the legacy revenue”).

An Ounce of Prevention...

When sourcing SIP trunking, consider the following up front:

- Interoperability between your specific IP voice and related equipment (gateways, Session Border Controllers, IP-PBXs, etc.) and the carrier’s network (this has improved greatly in recent years),
- Bandwidth requirements at both the data center and remote WAN locations (a function of traffic volumes, percent “on-net” vs PSTN-bound, codec choice for voice compression, supplier ability to aggregate bandwidth across the network in a “burstable” model, etc.),
- Legacy voice usage elements and how they are presented in the SIP future state pricing request (interstate vs. intrastate, features, etc.). Know your traffic!
- Availability of the supplier’s SIP trunking service across your corporate footprint. The request for proposals is a chance to pin this down by providing a list of your U.S. locations and asking each vendor to identify locations where its SIP service is not available (and the timeline for it to be available),
- Any potential problems with existing contracts (such as early termination penalties, minimum revenue commitments/potential shortfall impacts, etc.), how they might impact the procurement and deployment timelines, and possible plans to mitigate.

Big Savings Potential

Lest we sound overly negative, when properly done (which can be over time—you don’t need to flash cut all locations to SIP simultaneously), SIP transitions can produce substantial savings and free you from the tyranny of your least favorite telecom vendors/contracts (that would be the ILECs). SIP is also an area where some of the second tier vendors—notably Level 3, CenturyLink, and XO—have good offerings at good prices.

Get Help

Finally (marketing alert) for the reasons listed above and a few others, this is likely an area where the first time you do it you want some outside help. As they say to first year medical residents: watch one; do one; kill one.

Good luck in your SIP trunking sourcing efforts. 🍀

Three Things You Really, REALLY Need to Know Before You Tackle SIP Trunking

Hank Levine

What you need to know before you junk your business lines and PRIs in favor of the best thing to hit voice since direct dialing

In the late 90's it was Frame Relay; five years ago it was MPLS. Now SIP trunking is all the rage in corporate telecom circles. There are two good reasons for that. First, it saves large users a lot of money in an area (voice, and particularly local voice) where savings have been hard to get in recent years. Second, the incumbent local exchange carriers are the carriers that every enterprise loves to hate (high rates, lousy terms, arrogant attitude) and SIP trunking finally gives you a way to show them the door. Like anything new, SIP trunking has "quirks" and "gotchas". Here are three of them.

1 SIP trunks aren't trunks. IP (for "Session Initiation Protocol") is a set of standards for setting up calls and providing voice features over a packetized—typically MPLS—network. What we call a SIP "trunk" is actually a reservation of capacity sufficient to connect and carry a number of simultaneous voice calls over a data circuit. Each call occupies a logical circuit, more accurately described as a "concurrent call path".

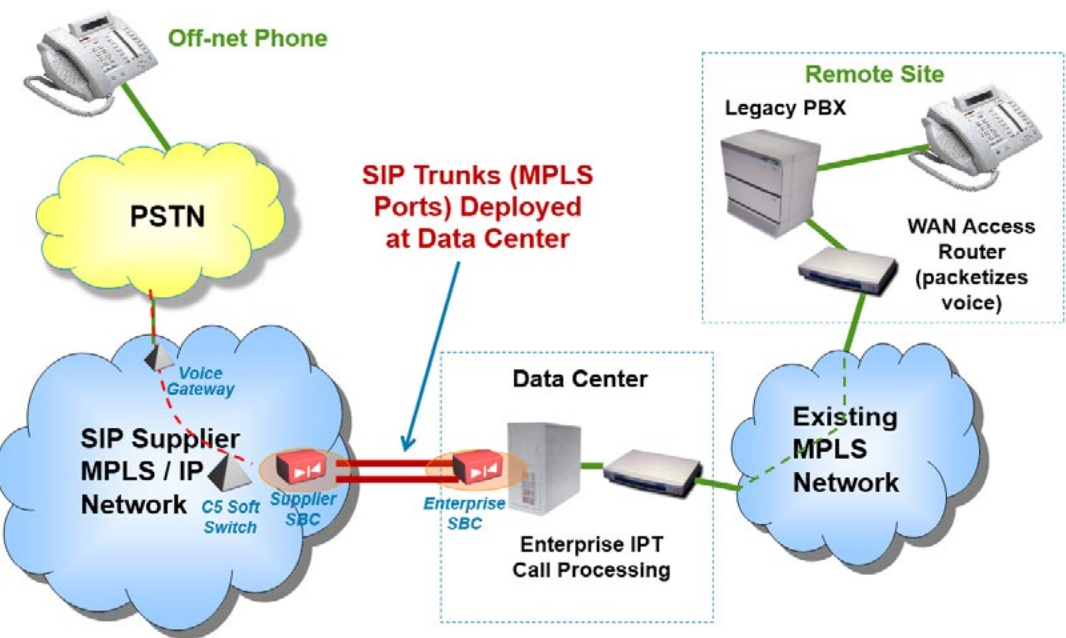
2 Now that SIP has moved from bleeding edge to off-the-shelf, it turns out that the biggest area of uncertainty and cost risk is the last 200 yards—**fiber connections to buildings**. Because SIP runs over the enterprise data network, at locations with a lot of users moving to SIP adds to the pressure from growing bandwidth demand for fatter access pipes. In the current environment that often means moving from a T-1 or T-2 to Ethernet access even at small and medium-sized locations. And that often requires optical fiber. Ethernet over Copper is not what it's cracked up to be. It rides over bonded copper pairs and degrades with distance; without repeaters, and assuming a bandwidth need of 10Mbps or greater, you need to be a mile or less from a central office if you have Cat-3 voice grade copper. So bringing Ethernet access to a building that isn't already "pierced" by optical fiber typically requires spe-

cial construction. And special construction costs a lot. The biggest complaint that we hear from large enterprises moving to SIP is that they negotiated what they thought was a big allowance to cover their anticipated special construction needs only to find that they were spending twice as much (and it was taking twice as long) as they anticipated.

3 As many times as we've talked to smart users about SIP, it always seems that half the audience doesn't realize that although you buy PBX trunks at every location, **you don't buy SIP trunks at every location**. Actually, using the most common configuration you only buy SIP trunks from one or two locations, typically data centers. The diagram on this page is from the SIP 101 session at the recent CCMI conference on negotiating telecom agreements, led by TC2's Jack Deal and David Lee, and LB3's Deb Boehling.

The only SIP trunks in this diagram are the red lines (which, to be fair, may be 100's of megabits each) that connect the customer's data centers to the SIP supplier so that calls can be sent to, and received from, the public switched network. If you are in NY, your data center is in Colorado, and you call the pizza joint down the block, the call will be routed over your data network to Colorado, where it will be sent to your SIP provider — over a SIP trunk — who will terminate it to the pizza joint. 1800 mile round trip; typically no usage charge for the 'local' call. That's a key reason why SIP saves a lot of money and why, when you use SIP trunks, you can ditch the local phone company. And—oh, by the way—you also don't have to get your SIP trunks from the carrier that provides your MPLS services.

If you are an enterprise customer who isn't already migrated or migrating to SIP, you probably will be soon. And now there are at least three things that won't trip you up. ☎



SIP Trunking is Driving Users' Ability to Consolidate Vendors in Europe

Ben Fox

Enterprises with multiple locations spread across Europe commonly purchase traditional fixed voice transport services (i.e., voice circuits, usage and DID/DDI services) from a range of local PTTs and in-country vendors. Even though a number of vendors exist that can offer such services across most countries in Europe, it is quite rare for enterprises to have consolidated these voice transport services with one vendor (or even a couple of vendors). And even when a single vendor has been engaged across Europe, the local PTTs still tend to linger on, e.g., keeping DDI services and local exchange lines.

Strategically procuring voice transport services across Europe drives significant cost savings, not least because in-country PTTs are almost always the most expensive vendor option. But various barriers and hurdles have historically discouraged pan-European vendor consolidation and strategic sourcing in this area:

- Obtaining the existing spend, pricing, contract and volume/inventory data tends to be an arduous task across so many incumbent vendors.
- Budgets for voice transport services are often held locally, so decisions are typically made locally too.
- Awareness of the total spend is often limited (not least where it's spread over a large number of vendors), and thus the services often fall below the radar during cost reduction initiatives.
- Voice transport services typically have lower management visibility than core data network and mobility services.
- Services and spend are sometimes combined with other in-

country services, such as mobile, and it can be awkward to de-couple and extract them.

- Incumbent PTT vendors often seek (and obtain) pricing/contract arrangements that make it hard to move away without incurring additional costs. These include auto-renewing contracts, prices that increase as volumes decrease, and discounts that can be "clawed back" if spend commitments are not met.

But SIP trunking has now become a force in the effort to overcome these barriers. A centralized SIP trunking architecture that consolidates all European PSTN access (ingress and egress) at your primary data centers forces vendor consolidation and reduces individual in-country PTTs to providing limited local site services (e.g., exchange lines for calling emergency services or out-of-band access to network devices).

The vendors' SIP coverage is continuously growing and the larger European and global SIP trunking vendors can migrate your voice services (including in-country DDI numbers) to their networks for most major European countries, supporting the majority of a large user's voice transport needs. This delivers a double dose of savings — not only the transformational savings from decommissioning expensive local voice trunks and replacing them with consolidated access at data centers, but also reduced outbound calling rates via the SIP trunks. Those rates are far lower than the in-country PTT voice rates that they displace. ☎

What the FCC's "Open Internet" Ruling REALLY Means for the Enterprise

Colleen Boothby

For enterprise customers, there's no question that the FCC's net neutrality decision tomorrow will be historic. (And no, not just because it's one of the very, very few times that the lobbying juggernaut of cable+telco hit a wall.) The decision is a long overdue, formal acknowledgement that the Internet is the dominant means of plain old communications in today's world. And it is way past time for Congressional leaders and D.C. policy-makers to admit that. Much of the discussion about the FCC's proceeding has focused on two big picture, public policy issues: Title II reclassification and forbearance. But for enterprise customers, the Order is also supposed to address two practical matters with high dollar impacts: terminating charges and universal service.

Reclassification

The Title II reclassification issue is whether the FCC should classify Internet access as a "telecommunications" service that is subject to the FCC's jurisdiction or an "information" service that is not. Classification as telecommunications does not determine whether the service will be regulated. See the forbearance discussion below. So in many ways, this issue really matters only to telecom nerds and the companies providing the service. Nevertheless, the issue has become quite the political hot potato even though it is really a factual, not policy, question — does broadband Internet access meet the definition of "telecommunications" service in the Communications Act? Is it the transmission of content chosen by a user between points specified by the user with no net change in format?

Back when the FCC first classified Internet access as an information service, we were still in a dial-up world. There was no disputing that the call to an AOL server farm was regulateable "telecommunications" and AOL's value add was a collection of un-regulatable information services. Today, there should be no disputing that the dumb pipe connecting my device to an Internet server is similarly "telecommunications."

That simple factual question got obfuscated by a lot of provider self-interest and craven political posturing. So it looks like the Wheeler FCC is finally do-

ing the intellectually honest thing — acknowledging that broadband Internet falls within the definition of a telecom service in the Communications Act and then using other tools in the Act, like forbearance, to deregulate wherever marketplace forces are strong enough to protect consumers. We can argue about whether the FCC gets the competitive analysis right but arguing that today's Internet access is not telecommunications is just denying reality.

Forbearance

The Communications Act authorizes the FCC to "forbear" from enforcing any provision of the Act if enforcing it is not necessary. When long distance markets are robustly competitive, for example, FCC price regulation is not necessary — competitive alternatives let consumers vote with their feet if a price is too high or service quality is too low. That is why the FCC forbears from enforcing the Act's requirement that long distance companies file tariffs. Tomorrow's net neutrality Order will do a lot of forbearing; it will "apply fewer sections of Title II than have applied to mobile voice networks for over twenty years," according to an FCC fact sheet about the draft Order. The question of which provisions should get forbearance and which should be enforced was another one that only telecom nerds and the affected companies could love. But it is also the question with the most impact on providers and will no doubt receive a lot of attention in the inevitable appeals of this order to the courts.

Terminating charges

The big, giant, huge win for enterprise customers in this order is supposed to be a "bright line" prohibition against double-dipping by Internet service providers. Oh we've all called it "blocking," "throttling," and "paid prioritization." But once again, let's call it what it really is: provider attempts to double-dip by charging their subscribers for connections and then charging companies trying to communicate with subscribers for the very same connection. Because from the enterprise customer perspective, here's the scoop. Consumers pay for their connections to the Internet. Businesses pay for their connections to the Internet. Why

should businesses have to pay somebody else's provider for somebody else's connection? Why should the consumer's provider be able to collect fees from its customer and from the businesses trying to communicate with that customer for the very same connection? Here's why: because businesses would have no choice but to pay once the subscriber signs up with a provider if they want their traffic to get through to that subscriber. Provider attempts to exploit their "terminating access monopoly" with charges to terminating content providers is exactly the kind of market failure that the FCC is supposed to regulate.

The FCC Order will ban providers from blocking traffic, degrading traffic on the basis of content, and charging for preferential treatment (the so-called "fast lanes"). This is a huge win for enterprise customers. Take that bull's eye off your chest. The cable companies and telephone companies can't force you to pay them as a condition of letting your traffic (your web site, your warranty information, online banking, sales channel, insurance claim processing, product specifications, etc.) get through to their subscribers.

Universal Service

The disappointing, infuriating, politically-driven punt in this order (thanks Congress!) is the FCC's apparent decision to forbear from requiring providers to contribute to the Universal Service Fund ("USF").

Enterprise customers currently pay a very expensive USF surcharge on every telecommunications service. That's because the FCC's rules require providers of telecommunications to pay into the USF based on their revenues from the sale of telecommunications. Pretty much all of the providers pass that contribution through to pretty much all of their enterprise customers in the form of a surcharge which has been hovering at around 17%. Seventeen percent! That's higher than nearly every other tax, fee, or surcharge applied to telecommunications.

The size of the surcharge depends on how big the base of telecommunications revenues is. But that base has been shrinking for years as residential

Continued on page 24

Negotiating Private Cloud Transactions: Building the Foundations of Best in Class Deals

Ben Fox and Marc Lindsey

Part I

Introduction and Pricing

Private cloud installations are an increasingly common component of enterprises' data center strategies. They complement traditional data center infrastructure and public cloud services (such as Amazon Web Services, Google Compute Engine, and Microsoft Azure).

Private clouds are essentially a do-it-yourself version of public cloud services. To build a private cloud, an enterprise purchases and then (on its own or through a third party) installs, hosts and operates dedicated equipment, network facilities and software to provide inter-

nal business units and customers with elastic, virtualized and pooled compute, network and storage resources; metered resource usage; and automated self-service provisioning capabilities. These are some of the key benefits of public clouds, and they come without many of the risks some fear when contemplating the multi-tenant shared infrastructures of public clouds.

Attracted by the agility, operational flexibility, efficiencies and scalability of public cloud computing, wary enterprises resist a full leap to public cloud

computing and turn to private clouds instead. The reasons cited for this decision include:

- IT Security requirements that are incompatible with a public cloud;
- Regulatory restrictions associated with certain types of data/applications;
- Custom requirements that do not map well to public cloud services;
- A need for service and security management/control unavailable to public cloud customers; and
- General nervousness associated

with exposing sensitive company or customer data to public cloud services.

Some enterprises build their own private cloud environments. Many, however, turn to system integrators and other IT service providers to install, host and manage dedicated private cloud infrastructures, blurring the lines between public and private clouds. When private clouds are built using third party services, the resulting equipment acquisition, software licensing, hosting and professional services arrangements can be complex. Success depends on thoughtful planning. This article examines some of the important deal points enterprises should consider when developing and executing successful sourcing and vendor negotiation strategies for hosted private cloud transactions.

MANAGE YOUR EXPECTATIONS

Private cloud arrangements cannot replicate the flexibility, economic efficiencies, and (seemingly) limitless scalability of public cloud services. Public cloud computing is fundamentally a utility service that can rapidly and flexibly absorb increases and decreases in the required workload, with corresponding increases and decreases in metered service charges. In contrast, a private cloud is a purpose-built, virtual environment dedicated to a particular customer. It is not a utility service. In a hosted private cloud, the provider purchases and deploys dedicated equipment and software on your behalf. If workloads drop, the service provider will still expect to be reimbursed for the outlay (plus a margin). Unlike a public cloud scenario, the charges in a hosted private cloud may not fall with decreasing workloads; indeed, the charges may be based on your projected maximum workload.

Enterprise customers often assume that service providers can provide private clouds on a far more flexible basis than is actually feasible. Private cloud services are not "productized" like many services. Service providers, for example, do not routinely re-deploy or reuse infrastructure for different customers. Private cloud deals (including scope and pricing) are assembled on a semi-custom basis, often drawing from a list of equipment

and software supported by the provider. The resulting hosted environment is customized to each individual customer instance, as is the pricing.

The dedicated nature of hosted private clouds means that the provider expects to recoup all of its outlay from each customer. The workload and cost flexibility that public cloud services offer are, therefore, incredibly difficult to replicate in private cloud deals.

Nevertheless, private cloud arrangements fulfill certain needs that public cloud services cannot. Understanding (and being sympathetic to) the areas where service providers have limited flexibility and, equally, knowing where the service providers can and should be flexible, will enable you to negotiate a best in class private cloud deal.

COMPETITION IS CRUCIAL

The most important factor in achieving a best in class deal is creating competition for your business. And the best way to drive competition is via a competitive procurement process involving multiple bidders.

Most of the deal and contract concepts discussed in this article should be tackled early in the competitive process (e.g., by including them in a Request for Proposal document issued to multiple service providers). Negotiating terms while you have multiple providers fighting for your business, rather than waiting until you've selected your preferred service provider and are close to the end of the negotiation process, will help you achieve the best possible outcomes.

PRICING AND COMMITMENT GUIDANCE

The key to a best in class deal is competitive pricing.

Pricing for private cloud deals is highly customized, but typically covers the following:

- Amortized capital costs, e.g., for hardware and software licenses, plus installation;
- Recurring costs for space, power and HVAC;
- Hardware and software maintenance;
- Managed services (ongoing monitoring and management of the pri-

vate cloud infrastructure);

- Bandwidth/connectivity costs; and
- One-time costs for implementation, and service changes over the life of the contract.

Service providers often combine cost components (e.g., rolling capital costs and space, power and HVAC costs) into a bundled recurring charge that applies on a per server/appliance basis. Understanding the relationship between the underlying costs and the service provider's pricing construct will help you to negotiate improved pricing and compare pricing across multiple proposals.

It is best practice to keep the pricing structure aligned with the underlying cost base incurred by the service provider. You may be tempted to ask a service provider to offer pricing for private cloud services that more closely resembles public cloud pricing (e.g., per user, per instance, or per workload, as appropriate). But these constructs are rarely aligned with the service provider's underlying costs and could create scenarios where the service provider is unable to recover its capital costs. To eliminate the risk of these scenarios, providers simply increase the offered pricing. Competitive pricing is closely aligned to the service providers' costs for the core pricing elements referenced above.

Negotiate charges for service changes up front

Providers will almost always try to levy additional one-time fees for changes to the private cloud infrastructure (e.g., to modify software configurations). One-time charges also typically apply when adding services or new resources, or increasing capacity of existing resources after the initial implementation. And don't be surprised to see higher charges for changes performed on-site, rather than remotely, even though the infrastructure is hosted at the service provider's facility.

The starting point for all such charges can be very high, so it is crucial to negotiate rates rather than discover how high the charges are the first time you request a change. It is also worth negotiating fixed prices for a menu of changes by type instead of (or in addition to) hourly rates.

Avoid overpaying for the supplier's upfront capital outlay

From the customer's perspective, it is perfectly reasonable to assume that the charges for the services will decrease once the service provider has recouped its upfront outlays for software, hardware and installation. But service providers often object to the concept, and do not proactively offer pricing that declines after the upfront investment has been recovered. It is therefore critical to lock in an appropriate price adjustment by including in the pricing schedule the date or event that triggers the reduction, and the actual amount of that reduction. Avoid hollow commitments to renegotiate pricing later in the contract term, at which point you'll have limited leverage to negotiate a fair adjustment.

Manage lock-in and commitments

Commitments and vendor lock-in arise in various ways in private cloud deals. Overall spend and volume commitments are readily identifiable and negotiable. But you need to address them early in the negotiations; don't wait to address them when you are drafting or marking up the contract.

Commitments can arise in other ways, too. In-service commitments, for example, arise when individual service components, once installed/activated, cannot be de-installed or deactivated until a certain period of time has elapsed. If the customer de-installs or deactivates a service before the end of its minimum in-service period, the customer will be required to pay service charges for the remainder of the period, even if it does not use the service.

For charges associated with the service provider's initial capital outlay, it is reasonable to accept a commitment to pay the associated charges for the time period over which the service provider has amortized the capital costs. But for most other types of services charges, minimum in-service periods should not apply (e.g., for management services, maintenance, space/power/HVAC charges, bandwidth/connectivity charges).

Bundled charges are also a source of concern. If amortized capital costs are bundled with managed services charges

and space/power/HVAC charges, the financial penalties for not fulfilling the minimum in-service period should not be 100% of the bundled charges. They should be a lower percentage of the bundled charge, ideally related specifically to the provider's upfront capital outlays.

It seems obvious, but including an express right to de-install services is a money saver. Service providers often structure contracts so that they commit customers to retain (or at least pay for) service components for the entirety of the contract term. This is another form of hidden lock-in/commitment.

Negotiating the notice period for de-installations also helps contain costs, and aligns the customer's total charges with its resource utilization. Ideally, charges for service components will cease as soon as you provide a de-install order (regardless of when the service provider gets around to physically de-installing the service component), but service providers typically ask for 30 to 60 days' notice. This delay is antithetical to the benefit of cloud computing, where most resources ought to be provisioned and de-commissioned using automated tools. Long de-installation dates should be limited to hardware.

Service providers typically amortize their upfront capital costs over the initial contract term in hosted private cloud deals. They do this in order to keep recurring charges as low as possible. It creates added complexity when service volumes are expected to grow over the course of the contract, because the time-in-service commitment for elements added part way through the contract term will extend beyond the end of the initial term. This decreases your flexibility and ability to switch providers at the end of the term. Developing and negotiating options that tackle contract term-extending in-service commitments should be front and center during the negotiations.

Minimize the service provider's rights to change prices

In general, the provider should agree to fix the rates and charges through stabilized unit rates or fixed monthly recurring charges for in-scope services for the term of the contract. Appropriate and

reasonable exceptions to "fixed" pricing may relate to labor costs incurred for services provided on a time and material basis, but even this is highly negotiable. Power-related charges are a common, yet difficult, exception to stable private cloud pricing. As with traditional co-location deals, those hosting private clouds usually seek to reserve the right to increase charges if electricity costs increase. It is generally unrealistic to completely eliminate such rights. A better course is to focus on managing their impact:

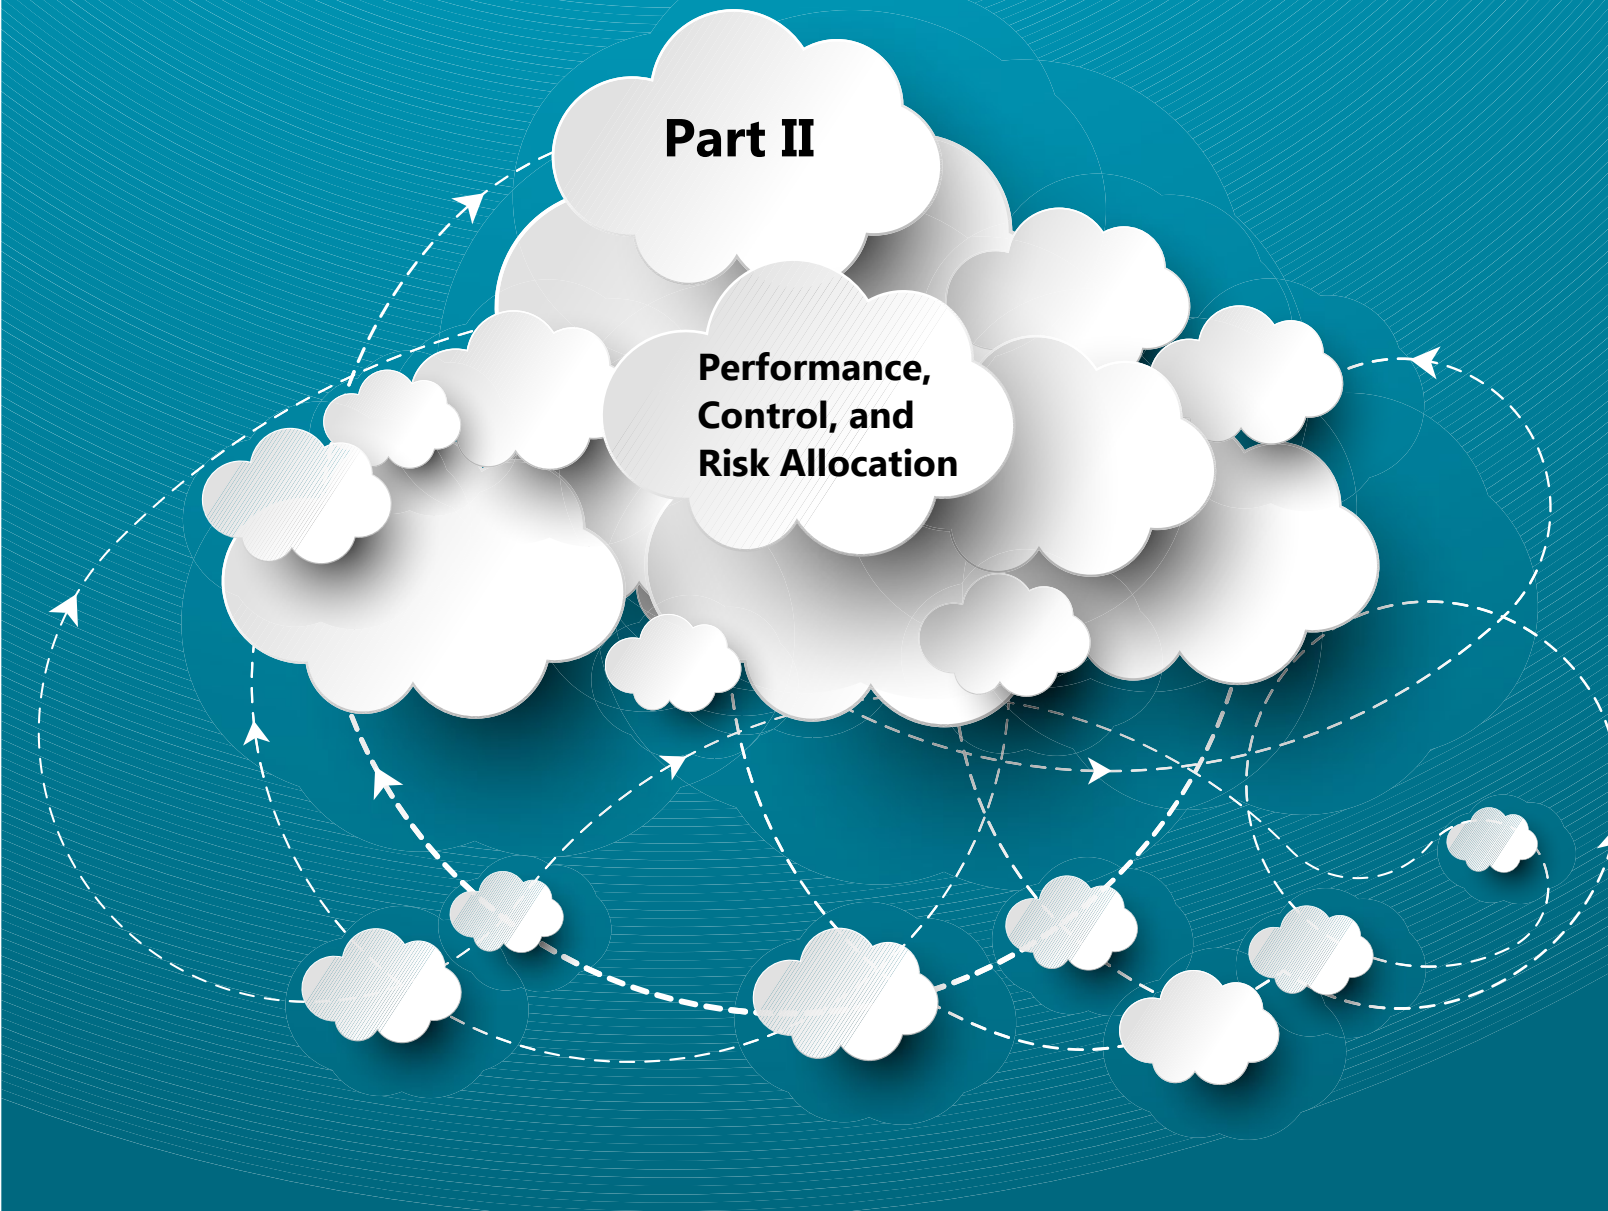
- Any price increases should be in direct proportion to the increase in energy prices, and should only affect the portion of individual charge elements that are related to the cost of power.
- Prices should go down as well as up.
- Prices should only increase in relation to a material change in supply costs.
- Service providers should only be allowed to increase prices periodically (e.g., no more than once per year).
- The service provider should bear some burden of showing that its costs have increased, and by how much.
- Although it does not typically translate into specific pricing parameters, customers should require their providers to disclose the power usage effectiveness ("PUE") at the data center where the private cloud is hosted before inking the deal, and then require the providers to regularly report PUE. Comparing this figure to industry standards can help inform power pricing negotiations.

Nail down trial period pricing and exit costs

Most private cloud pilots or trial periods will still require the service provider to fund significant upfront costs, particularly on hardware and software. It is unrealistic to expect to be able to exit a contract after a pilot or trial period at no cost, particularly if the service provider is not at fault (e.g., if you conclude that your planned workloads are simply not suitable for a virtualized environment). Pilots and trial periods can be negotiated, but expect there to be shared risk and a shared cost burden. ☎

Negotiating Private Cloud Transactions: Building the Foundations of Best in Class Deals

Ben Fox and Marc Lindsey



FINE TUNING SCOPE AND PERFORMANCE

Carefully document the scope of the services

Whenever dealing with outsourced or managed services, it is essential to carefully document the services to be performed and the responsibilities to be fulfilled by the provider. Lack of clarity on where in-scope (no additional charge) ends and out-of-scope (additional charges apply) work begins can

unravel an otherwise good deal. Avoid misunderstandings and additional costs by clearly defining the following at the start of the process:

- Will additional charges apply for routine changes to the services or are such changes included in the managed services fees?
- What end user support, if any, is included?
- Who is responsible for capacity management?
- Where is the dividing line between the provider's management re-

sponsibility and the applications that the customer will manage and install?

It is also useful to negotiate specific terms regarding the ability to grow the service, specifying an upper service volume/capacity that the service provider commits to make available if requested by the customer. Negotiating room to grow is particularly important where a private cloud environment is located in a single data center, rather than being spread over multiple data centers.

Negotiate robust service level guarantees

The providers' standard SLAs are notoriously weak — designed to look good on cursory review while avoiding liability for meaningful service credits under almost any circumstances. Experienced customers negotiate custom service level guarantees that are aligned with the importance of the services at issue to their businesses, and that suitably incent providers to do all they can to maintain service availability and operability. Key service level requirements include:

- Overall availability of the private cloud services. Overall availability of the services, rather than just availability for individual service components, is important. For example, if users cannot access any services whatsoever because of a cross connect issue, then a service credit based on a percentage of a small monthly cross connect charge is unlikely to be satisfactory. Availability of individual components, or time to restore individual components, should complement overall availability.
- Initial implementation guarantees. Service credits must be payable in the event that the implementation of the services is not completed within agreed timeframes.
- Time to complete changes or add new resources. There should be guaranteed timeframes within which the service provider will complete standard service changes, upgrades or service additions (e.g., adding users or incremental infrastructure components). For true private clouds, many routine resource provisioning requests should be automated, and occur rapidly.
- Service availability during maintenance periods. The service provider's maintenance activities should never mean that services are unavailable. A robust private cloud service will have sufficient redundancy that maintenance can be performed without actually taking down the production environment.

COMMAND, CONTROL AND ACCOUNTABILITY

Maintain control

In private clouds, enterprise customers lose most of the pay-as-you-go cost flexibility available in the public cloud context. Inevitably, as discussed in Part I, building a private cloud or having one built for you means you must pay for the licenses, hardware and implementation services required to deploy the cloud, whether in the form of upfront purchase with capital, non-recurring implementation charges and/or fixed-duration recurring charges. In exchange for this cost disadvantage, private cloud customers receive dedicated environments, which means they should have considerably greater control. Indeed, greater control is often a principal driver for selecting a private instead of a public cloud.

An important element of control is the transparency of critical IT service management processes. In private clouds, the customer's change and security management requirements should not presumptively yield to the provider's standard processes.

Thus, in public clouds, customers are beholden to the provider's pace of change and upgrade cycles, and must accept much less visibility and participation in service management of the provider's multi-tenant, shared infrastructure. Since private cloud customers pay for the privilege, they get to dictate critical aspects of change management. For example, you should prohibit your provider from introducing material and adverse changes into the private cloud environment; require all testing and upgrades to be coordinated with you; and preserve the right to reject upgrades or changes you may find undesirable.

Here are some other important private cloud contract controls:

- Your provider should not be allowed to relocate data centers or move the environment within the selected data center without your consent;
- Because private clouds are ordinarily purpose-built for customers and usually come with substantial financial commitments, you should get a commitment from your service pro-

vider to support the environment for the term. Any provisions that give the provider the right to discontinue service (except in the case of your uncured material breach) puts your investment at risk;

- You should have full access to perform vulnerability and penetration testing, and to audit your private cloud environment for security, data privacy and regulatory compliance purposes; and
- The contract should draw clear lines of responsibility for, and regulate use of, the test environment. It should also document policy-enforced workflows for moving items from test to the production private cloud environment.

Fairly allocate responsibility and liability for security incidents

A third-party provider that operates, hosts, and manages private clouds for an enterprise customer has actual control over the physical—and much of the logical—security of the environment. But the customer also bears responsibility for good security practices with respect to its access, and use of the private cloud, including (ordinarily) management of the applications themselves. Although it is often the most critical factor in moving to private clouds, the contractual allocation of responsibility for security management between provider and customer is often overlooked.

Drawing these lines in contract documents before work starts can be difficult, but it is essential to effective compliance management. It exposes potential misunderstandings, giving the parties a chance to identify and remedy them before entering into contract.

Working out the allocation of operational responsibility is hard work for the technical and security-focused deal teams, but it is not usually controversial. That said, getting providers to accept financial responsibility for damages caused by their failure to perform the security and privacy-related obligations assigned to them can be very difficult. Providers typically seek to avoid any liability for loss, corruption, or im-

proper disclosure of customer data and applications, even when caused by the provider's contract breach, negligence or misconduct. Enterprise customers, on the other hand, expect these sorts of liabilities to be uncapped. Both providers and customers have been willing to walk away from transactions over the issue.

There is no single solution to this problem, and your risk management team should develop a clear sense of the risks your company is willing to accept (with fallback and bottom line positions), ideally before commencing negotiations with any specific provider. At the very least, you should only do business with providers that are willing to accept reasonable financial responsibility for losses resulting from their security failures. To close deals with risk-focused but otherwise reasonable suppliers, you also have to be comfortable accepting maximum exposure limits for them, subject to uncomfortably narrow exceptions, e.g., uncapped liability for fraud, willful misconduct, violations of law, and certain types of remediation costs, with other types of damages capped at an agreed upon multiple of the contract's annual value.

CONTINUITY AND BUSINESS DISRUPTION

Private clouds are often preferred to public clouds when sensitive, highly regulated or other critical customer data is involved. This ordinarily means the applications destined for private clouds support mission critical or competitively significant business processes. For these critical applications, designing the environment for high availability operations and tight security is the best protection. In hosted private cloud transactions, threats to business continuity are therefore key contractual considerations.

Some suppliers like to reserve liberal rights to suspend services when they, in their sole discretion, determine suspension is appropriate. Efforts of this kind should be identified and removed. In a private cloud context, a supplier's right to suspend the customer's access to, or full use of, its private cloud environment should only arise when the

customer fails to cure a material breach of contract or engages in malicious conduct that materially threatens the security, stability or normal operation of the vendor's infrastructure or the services provided to other customers. It is important to note that contractual suspension rights are distinct from security incident management. To respond to and remediate security incidents, suppliers may need to temporarily restrict access until the situation is resolved.

Provider termination rights, like suspension rights, should be narrowly constructed. Termination should only be permitted for material breach of the negotiated terms and conditions where the breach remains uncured after an appropriate cure period.

Non-payment by customers as a termination event can be particularly tricky. Providers rarely accept materiality as a precondition on non-payment related terminations. Customers, partly in recognition that provider billing practices are notoriously error prone, want some leeway for payment delays. The typical compromise allows providers to terminate for non-payment, subject to prescriptive notice periods and the customer's right to withhold payment of any amount it disputes in good faith.

When the contract expires or is terminated for any reason (including the customer's breach), suppliers should be contractually bound to render migration assistance services that allow the customer to transfer the private cloud environment and related services to another provider or bring them in-house in an orderly fashion; migration services include continued provision of the in-scope services for a defined period (e.g., 9 to 12 months depending on the complexity of the environment), along with additional handoff activities (e.g., documenting the elements and configuration of the environment, transferring software licenses and hardware assets using a pre-negotiated pricing formula).

When the termination is the result of the customer's non-payment of undisputed amounts, it is reasonable for the provider to condition the migration services on the customer's pre-payment of a portion of the charges for service to

be provided over the migration period.

PRIVATE CLOUD SERVICES ARE GROWING, AND WILL BE WITH US FOR SOME TIME

As public cloud services mature and evolve, enterprise customers will become more comfortable using them for sensitive and mission critical services and data. We expect that to happen over a relatively long horizon. In the meantime, enterprises will continue to deploy private cloud installations over the next several years to reap some of the advantages of cloud computing innovations.

Hosted private cloud services represent a significant proportion of private cloud installations, and are expected to lead enterprises to greater adoption of public cloud services. Many vendors offering both private and public cloud services are highly motivated to win hosted private deals to position themselves for future public services. The cloud providers' eagerness to sell cloud services, along with fierce competition amongst private and public cloud providers, allow enterprise customers to negotiate solid agreements and obtain market leading services if they understand and embrace evolving cloud deal best practices. ☁



Why Ethernet Access is a Critical Part of Enterprise Wireline Networks

Hank Levine and Jack Deal

In today’s Brave New World of enterprise telecom, MPLS has supplanted private lines, frame relay and ATM as the backbone of data networks. At the same time, SIP trunking is well on its way to replacing TDM-based PRI’s for voice. By the end of this decade—at the latest—MPLS and SIP will be the near-universal model for enterprise wireline communications.

Sounds like a great plan. And it is.

The fly in the ointment—or, if you prefer, the bug in the code—is dedicated access, the pipe your voice and data traffic rides from your corporate locations to the nearest network node. In other words, the last (and first) mile. Access is not trivial—traditionally it accounts for 30-40% of your total network costs and for decades it has been the thinnest connection (hence the bottleneck) of corporate networks. And since it’s the last bastion of the ILEC’s near-monopoly when CLEC alternatives aren’t available, it’s notable for high prices and lousy terms.

Until recently, the workhorse of dedicated access was DS-1 [1.5 mbps] special access or, for really large locations and data centers, its big brother, DS-3 [45 mbps]. Separate lines were procured for voice and data (a PRI is a DS-1 divided into 23 voice channels).

But today, DS-1’s don’t cut it anymore, DS-3’s are expensive, and carriers are making noises about discontinuing both. Voice is now an “app” that increases the load on special access lines dedicated to data traffic. And even without voice, who wants 1.5 mbps speeds when FiOS or Comcast gives you 10-30X that at home? Finally, bandwidth demand is soaring as employees use more and more cloud based apps.

ETHERNET TO THE RESCUE

The good news is there is no real debate about the replacement for TDM access, and we don’t have to wait for it to mature because it has been deployed for decades. It’s Ethernet.

Ethernet grew out of ALOHAnet, which was developed 40+ years ago in (where else) Hawaii. Over time, Ethernet beat out IBM’s token ring and became the standard for routing traffic over local area networks. About a decade ago it began to appear in carrier networks as Carrier, Metro, or Wide Area Ethernet.

Ethernet for dedicated access has many advantages, even beyond the fact that, unlike TDM, carriers aren’t trying to do away with it.

- It’s the native protocol on virtually all LANS, so using it in wide area networks is simple, and reduces the number of protocol translations required to traverse a network from end to end.
- It’s more granular and more easily scalable than TDM: going from 5 to 10 to 20 to 100 mbps in the Ethernet world doesn’t require bonding DS-1’s, and then converting to a 45 mbps DS-3, and then bonding those DS-3’s into 155 mbps OC-3 or higher circuits.
- It doesn’t require a lot of education—your IT personnel already know Ethernet.
- Ethernet customer premises equipment is much less expensive than TDM equipment.
- But most of all it is relatively cheap and getting cheaper (although it’s Layer 2, which means that it’s a telecom service subject to USF and associated surcharges). Today Ethernet access is less expensive than TDM for anything over about 5 mbps. For 10 mbps or higher—the new norm—the comparison isn’t even close.

THE PROBLEM

Is there a catch? Well, if all you need is 1.5 mbps, a DS-1 is still cheaper and more widely available than a 2 mbps Ethernet connection. But the real ‘gotchas’ of Ethernet access flow from the fact that it likes to run over optical fiber.

Before the second tier vendors descend on us like a pack of flying monkeys,¹ we hasten to add that Ethernet is independent of its medium and one can in fact buy Ethernet over Copper (EOC) and even Ethernet over Wireless (EOW). However, except as a backup and for cell site backhaul, EOW is more talked about than deployed. And EOC has limitations – it doesn’t work at locations more than 12,000 feet from a central office and degrades well short of that; it has bandwidth limits (vendors talk about 50 mbps but don’t count on anything more than 10 mbps); and it is subject to technical issues like throughput fluctuation and signal degradation. EOC can be useful, and it absolutely beats nothing, but twisted pair is not the ideal Ethernet medium.

So what’s the problem with fiber? In a word (OK, three words) it’s not ubiquitous. The most widely quoted estimate is

¹ In The Wizard of Oz, the Wicked Witch of the West never actually says “Fly my pretties, fly!” Of course, in Casablanca Rick never says “Play it again, Sam.” And no one in The Treasure of Sierra Madre ever says “We don’t need no stinking badges.” But we digress.

that optical fiber has been deployed to about 40% of U.S. buildings housing more than 20 employees, though that is increasing rapidly (in significant part because of demand for Ethernet access).² For smaller buildings the penetration is considerably lower.

That means a large user planning to deploy SIP trunking, video services or other bandwidth intensive applications across a region or nationwide (and/or contemplating a substantial upgrade of its data network bandwidth) will likely encounter what the telecom world calls special construction to reach a less than trivial portion of its key locations. And special construction has two problems:

- *Delay* – it can take months (3-6 months; longer if you’re unlucky) to actually bring optical fiber from the street (or several blocks away) into a building. It isn’t just the construction – it’s the time to deal with your landlord and secure building permits from the local authorities.
- *Cost* – special construction is not cheap. Typical quotes are \$100-125 per foot (!) to build a fiber lateral connection to an office building, which works out to around \$60K if you need to go 1/10th of a mile. Recently, Windstream filed a letter with the FCC that reported a special construction quote from AT&T for a single 10 mbps Ethernet circuit of \$99,685 plus \$53,125 in “revenue recovery charges.” And you can’t predict the financial impact of special construction in advance; each quote is prepared on an individual case basis (ICB) after you place an order.

Separately, these issues can be challenging. Together, they threaten to bust budgets and schedules in a way that enterprises do not expect and cannot tolerate.

HOW TO PROTECT YOURSELF

So what can you do? A lot, actually. Here are two general tips that will help (in this and lots of other areas):

1. Procure Ethernet access as part of the RFP for the services that will use it (most likely your MPLS network and SIP trunks). If you single source, you’re toast.
2. Remember that access is only one component of the total cost of what you’re buying—no one buys access to nowhere.³ Other costs include port and class of service (CoS) elements in an MPLS network solution. What really matters is your total cost of ownership.

² Vertical Systems Group - <http://www.verticalsystems.com/vsgpr/us-business-fiber-gap-narrows-in-2013/>. Used with permission of Vertical Systems Group, Inc. © 2014; all rights reserved.

³ Actually, Hank was once involved in a transaction where we did exactly that, but it’s too long a story to tell here.

And to deal specifically with optical fiber build-out issues:

1. In the RFP, list your principal sites and ask vendors which are already served—not just passed—by fiber to which they have access.
2. If you ask and your procurement is perceived as competitive, the vendors will give you an allowance for special construction. For a starting point, take what you think you’ll need and double it. You may have to agree to a lengthy commitment (overall and for each circuit) to get everything you need. Without that, be prepared to pay at least some costs up front.
3. Ethernet access pricing is not yet uniform across the country, though it’s moving in that direction. The most common approach is custom rates for distinct bandwidths that vary by “zone” or area code. Work to achieve a more standardized pricing model during negotiations, as that can substantially affect your costs when you need to add the next new location.
4. Finally, if you’re introducing (or greatly expanding) Ethernet access as part of your migration to SIP trunking, use the fact that SIP trunking can be rolled out to locations over time to plan your Ethernet access implementation. For example, if a building or office park isn’t served by fiber yet but there are firm plans to do that in a year, it may make sense to defer that location until the fiber is in place.

The bottom line is that Ethernet access is the perfect complement to MPLS and SIP and will become the “on ramp” of your wireline network over the next few years, if it isn’t already. It is a key source of savings and improved service, and since the technology is mature, deploying it is not a major risk. So avoid the pitfalls as you migrate, and enjoy the results. 📶

WHAT IT MEANS FOR 2015

The good news is that the problem described in this piece—special construction or the like for buildings not pierced by fiber—will diminish over time as more and more enterprise-level buildings are served by fiber. The less-good news is that though the problem will diminish, it is unlikely to disappear until after 2020, if then. So large users will need to continue planning for some special construction—with the attendant cost and delay—for years to come.

AT&T Issues Withdrawal of Service Matrix – Part 1

Deb Boehling and Janis Stephens

Early in December, AT&T added a “Withdrawal of Service Matrix” to the General Provisions of its Business Service Guide. The table identifies services that AT&T will cease to offer sometime in the (likely near) future, even to current customers of the service.

AT&T tries to soften the blow, but its efforts are ham-handed. First, it characterizes the new services as “more technologically advanced,” but doesn’t say whether AT&T or the customer or both will benefit from these enhancements. Second, it identifies “eligible replacement services” the customer may want to consider. Not surprisingly, AT&T mentions only AT&T replacement services, and doesn’t provide any assurances that they will perform as well or better than the services they replace, will cost no more than the replaced services, and/or will not require the customer to make expenditures to replace equipment that works well with the replaced services.

In truth, for a decade AT&T has allowed itself to discontinue offering services on 12 months written notice and service components typically on 120 days written notice, but it was silent about which ones — you had to wait for the notice to arrive. In that sense, the table helps customers by letting them know a bit more about what’s to come.

The table currently includes only one category: Local Exchange Services. Customers negotiating with AT&T over the last few months heard AT&T rumbling

about its plans to stop offering certain services and may be surprised at the brevity of the “Affected Services.” But AT&T then proceeds to list all services associated with its standard voice services — AT&T Business Network Services, AT&T

OneNet Services, and Prime Services — and of these only Prime is local.

What does this mean? If you use any of the specified AT&T Services, you might get a notice any time that your service is being withdrawn. Then you’ll have one year (or some other period if specified in your agreement with AT&T) to figure out what to do about it and to imple-

ment a new solution. Note that AT&T’s VTNS service is not included in the table, but it is likely to be added soon — AT&T has already announced that it will be sun-setting at the end of 2018.

And finally, what should you do? If you use any of the listed services, you need to start planning to replace them if you aren’t already doing that. AT&T helpfully lists their services that could replace the withdrawn services, but that doesn’t mean those are your only choices. You could turn to other suppliers — and should at least consider that option. After all, if you have to rip out old AT&T services, it’s no more work to replace them with an alternative service from one of AT&T’s competitors than to replace them with the services AT&T wants you to use.

Best advice — get moving! That year will be gone before you know it. 📞

“Many Enterprise customers understand the dangers of the carriers’ online terms and changes to them, but few have time to keep up with it all.

A recent move by AT&T demonstrates the importance of keeping abreast of the changes.”



AT&T Issues Withdrawal of Service Matrix – Part 2

Deb Boehling and Janis Stephens

In early December, we let you know about AT&T’s Withdrawal of Service Matrix. We expressed surprise that only POTS/TDM telephony voice services were included in AT&T’s December 3rd chart, and opined that “AT&T’s VTNS service . . . is likely to be added soon.” Sure enough, on December 14th, AT&T updated its withdrawal matrix, and VTNS is on the chopping block. Now we’re surprised in the other direction — the list of services that AT&T plans to withdraw is far longer and broader than we predicted.

AT&T now says that it plans to withdraw all non-Ethernet access channels (e.g., DS0, T-1, T-3, OC-3, OC-12, OC-48 connections you have to have to use AT&T’s designated “Eligible Replacement Services”), all non-Ethernet private lines (which AT&T calls “IOCs” and International “half-channel” or “full-channel”), and Ethernet private lines slower than 600 Mbps. And while suggested by AT&T’s December 3rd withdrawal of “[a]ll services associated with AT&T Business Network Service [and] AT&T One-Net Service,” AT&T’s December 14th and 19th updates leave no doubt that existing conferencing services and toll-free features will also be replaced.

Why do you care?

120 Days’ Notice Leads to Migration Migraines...

AT&T’s latest approach probably shortens how long you have to plan for the move. In its December 3rd announcement, AT&T listed the applicable Services/Service Components as “all services” under certain plans. On this basis, we urged you to start planning because 12 months will fly by. AT&T’s latest additions do not involve withdrawing “all” of a certain service; instead it lists service components (like DS0, T-1, T-3, OC-3, OC-12 and OC-48 access and private lines). That’s important because AT&T is required to provide only 120 days’ notice before discontinuing a service component, meaning a circuit, feature, or other chargeable part of an overall ser-

vice offering. So on as little as 120 days’ notice, AT&T can turn off certain of the voice, access, and metropolitan area and wide area network service components on which many enterprises depend. And if access must be replaced with Ethernet access, you will discover (if you haven’t already) that AT&T and its competitors all think installation of Ethernet access can take well over 120 days. In New Jersey, one carrier estimated installation of a single circuit would take more than a year.

If your enterprise uses non-Ethernet access or non-Ethernet private lines at U.S. sites (standard TDM or SONET access is common for existing networks, particularly those with a distributed architecture), every U.S. site must replace old TDM or SONET services with new Ethernet services. For many customers, non-Ethernet access services work well, provide sufficient bandwidth, and avoid implementation headaches and unknown costs associated with Ethernet access — e.g., special construction charges, long implementation timeframes, and delays meeting even those timeframes.

Over the last decade, we’ve seen most enterprises move to an MPLS-based WAN or virtual private line service to connect even major locations (e.g., data centers), but a few — particularly those in regulated industries or who are contractually obliged based on government security requirements — still use high bandwidth SONET private lines for connections between large sites. Moving these to Ethernet private lines with a minimum bandwidth of 600 Mbps will take time and raise the same concerns as moves to Ethernet access on a smaller scale. And if your enterprise uses DS0, DS1 or DS3 local private lines or OC-3 or OC-12 Ethernet metro private line services, AT&T doesn’t (yet) offer a replacement. You must upgrade to Ethernet OC-48 or OC-192 wavelength services — and more bandwidth means higher costs and burdensome service swaps.¹

If you are a large toll-free customer, pay particular attention to the inclusion of toll-free services in the matrix, and

note that toll-free services are still in the replacement services column — so maybe part of the service will be withdrawn and other parts will remain. Although it’s clear that AT&T is planning to discontinue TDM services in favor of IP, it’s unclear what AT&T intends to do with the features that large customers depend on to route toll-free calls. Since customers have been slower to implement IP toll-free than local and long distance outbound services, the withdrawal of toll-free services has the potential to affect many more customers, and create disruption of critical business functions, even if the toll-free features are not withdrawn. And you may well have to transition your toll-free services at the same time you’re dealing with the complications of replacing your TDM access with Ethernet.

... And Creates Business and Operational Risks

There’s more to the forced march than uncertain installation dates and migration risks. There may be purchase commitment risks, operational risks (e.g., will the SLAs for the replacement service be as good or better than those for the withdrawn service?), cost (e.g., will AT&T’s prices for the replacement service be as good as those of the service AT&T is withdrawing; can you buy the replacement service from a different vendor without increasing your costs?) and risks such as finger pointing between carriers. Review your contract with these things in mind. If it doesn’t include protections (e.g., reducing purchase requirements when AT&T “replaces” or “withdraws” a service; counting replacement services toward purchase requirements; requiring substantial notice before AT&T can discontinue a service component and more

before it withdraws an entire service), you’ll need to amend it. If that doesn’t work, you’ll need the FCC’s help.

What can you do?

Coming full circle, start planning now. Press your AT&T account team for details about when the services and service components your network relies on will be withdrawn. If you are in the midst of negotiations with AT&T, insist an “earliest withdrawal date” be included in the contract. If you are under contract with AT&T and ordering new service components, order the replacement services if possible so you don’t have more to convert later on. Ask other vendors to provide proposals for replacement services that could be implemented reasonably quickly and consistent with your requirements.

Finally, there may be hope for preventing (or at least affecting the timing of) the forced march. Under Section 214 of the Communications Act, a carrier may need Federal Communications Commission approval before it “discontinues, reduces or impairs service;” various states have their own requirements for intrastate services. It’s not clear how the statutes will affect AT&T’s plans or the timing/manner in which it withdraws a particular service. The FCC opened a formal proceeding late last year to address the consumer impacts of technology transitions in basic network services. Last month, as part of that investigation, the FCC issued guidance on the lead times and disclosures carriers must provide to customers and competitors before discontinuing a service. The FCC also requested input from interested parties on any changes to the discontinuance rules that may be necessary to protect consumers and competition during technology transitions.

Savvy enterprise customers should, at a minimum, monitor events at the FCC as they unfold to ensure that they have a solid understanding of their rights under the FCC’s rules and the negotiating leverage those rights may give them. Truly savvy enterprise customers also get involved in FCC proceedings to ensure that the FCC is aware of end user concerns and crafts its rules accordingly. If your company wants to get involved, you should consider membership in organizations like the Ad Hoc Telecommunications Users Committee that represent customers before the FCC. It’s the only way to make sure that the regulations work for customers, not just carriers. ☎

WHAT IT MEANS FOR 2015

The Withdrawal of Service Matrix and things like it are only going to grow over time. Which means that Enterprise customers need to pay attention to what’s on them and what that means for their procurement planning in terms of commitments, equipment investments and the like.

Verizon Rapid Delivery — Not SO Fast

Laura McDonald and Deb Boehling

Long delivery times are among customers’ biggest complaints about their telecom vendors, so Verizon’s “rapid delivery” automation platform is sure to catch the attention of enterprise users. It’s supposed to drive faster quoting, contracting, ordering and provisioning, and more accurate billing. What’s not to love?

A lot, actually, if you are an enterprise customer who understands that price is more than just a number and likes to negotiate appropriate terms when contracting for network services. Fast is good, but it’s not synonymous with value, especially when accompanied by hidden terms and higher fees for mission critical services. And there is little evidence that Verizon’s approach is actually yielding better, more quickly deployed, and more efficient services for customers.

VERIZON’S RAPID DELIVERY CONTRACTING MODEL

Every few years one of the vendors comes up with a “solution” to the age old enterprise complaint that their contracts are cumbersome and one-sided, and the delivery and implementation of new prices and services is slow. Unfortunately, the “solutions” are usually illusory. Verizon’s Rapid Delivery contracting model fits this trend in a way that’s similar to AT&T’s “Comprehensive Service Order Attachment,” introduced almost a decade ago. That initiative also led to confusion and customer dissatisfaction, and AT&T has since moved away from it.

Verizon recently started aggressively pushing its large enterprise customers – potential and existing – to a new contracting model it calls Verizon Rapid Delivery. The new model has the same elements as the old – (1) a Verizon form contract with (2) service-specific Attachments, both of which are subject to (3) Verizon standard web-based documents (such as its Service Publication and Price Guide and Acceptable Use Policies). The difference is in how it treats these elements.

Under the existing/old model, Verizon customers have a master service agreement with service-specific attachments that are just that – specific to a particular service, such as Internet Access, MPLS (which Verizon calls “Private IP Service”), web conferencing, security services, WAN optimization, etc. Under the new model, Verizon’s Rapid Delivery Attachment is set up so that it can be applied to all Verizon services that are “optimized” for its automation platform, and applies specifically to the services listed as “New Rapid Delivery Services”. Verizon no longer provides a copy of the service attachment for a specific service; the VRD Attachment incorporates by reference Verizon’s web-based terms.

The number of “optimized” services is large and growing, though Verizon has not yet included (and may never include) services that it wants to do away with, like TDM based long distance and web based telephone directory services. We’ve appended a list of “optimized services as of mid-July of 2014; it includes most of the services of greatest interest to enterprise customers,



such as S O N E T and Ethernet access, MPLS, Internet Access, managed WAN, Voice over IP, wavelength services, and Wi-Fi.

Why should you care and how is this approach different from what enterprise customers have experienced from Verizon (and other providers) in the past? Under the “old” approach sophisticated enterprise customers were able to avoid some of the worst parts of Verizon’s web-based terms and prevent many material and unhelpful changes to things like key functions and features. They did so by negotiating terms in their contracts to “trump” undesirable web-based terms and by copying desirable web-based terms into a service attachment or their master agreement to eliminate Verizon’s ability to change them. Sophisticated customers refused to allow non-negotiated terms “incorporated by reference” to undermine their negotiated contracts by specifying that negotiated terms would take precedence over non-negotiated, incorporated terms.

All of this matters to customers because under today’s standard order of priority, if there’s a conflict, service-specific terms in an attachment trump the master agreement, and the master agreement trumps the web-based terms. But Verizon’s Rapid Delivery muddies these waters: the Rapid Delivery terms take precedence over the master agreement and other attachments. And the Rapid Delivery terms include new, often less favorable terms and incorporate web based terms. They also ask customers to acknowledge that the Rapid Delivery terms and rates (and any modifications and incorporated terms) are binding.

Verizon’s approach draws into question which terms will apply. This is not a question any enterprise customer wants to leave unanswered. Negotiated terms and rates should always trump the online terms (and prices) that the vendor wrote and can change at will.

In short, it appears that Verizon is promising “rapid delivery” by eliminating negotiated terms and undermining existing ones. That is a trade-off few enterprise customers are willing to make for mission critical services.

¹ It’s not surprising that Frame Relay and ATM services are now included in AT&T’s Withdrawal of Service Matrix; for years they’ve been offered only on a “grandfathered basis” to existing customers. When AT&T pitched AT&T VPN services instead of Frame Relay and ATM many enterprises initially hesitated — what they had worked well and the benefits of moving were uncertain. (Many did not view MPLS’s “any-to-any connectivity” as a big benefit before voice services were easily packetized and offered near POTS quality.) AT&T tried to entice customers to move off of Frame Relay and ATM by increasing charges and emphasizing the savings available via AT&T VPN. Those enticements worked for many, but AT&T could not completely discontinue its Frame and ATM Services — both of which are on the matrix’s chopping block. AT&T wants to avoid repeating the past.

VERIZON THROTTLES SERVICES, SO CUSTOMERS SHOULD THROTTLE VRD

What New Customers Can Do:

New customers don’t have to accept Verizon’s “solution”. If you are a potential Verizon customer preparing to release an RFP, include specific prohibitions against a vendor using its order process or other sleights of hand to avoid complying with the agreement reached.

If you are a potential Verizon customer that released an RFP, but are still in discussions with a short list of vendors, use the other bidders contracting flexibility to negotiate a clause specifically excluding application of the VRD Attachment to any services provided under the agreement.

If you are a potential customer that released an RFP, but have already chosen Verizon as the winning bidder, ask for the clause described above and, if Verizon refuses, re-evaluate the bids.

What Existing Customers Can Do:

Over the past few months, the concern over the Verizon Rapid Delivery Attachment has grown, particularly when Verizon claims it will no longer process orders for the “optimized” services under a customer’s existing contract — a contract that appropriately reflected the user’s understanding of the business deal when it decided to purchase Verizon services. We have yet to hear why the automation platform could not be used, but with services remaining subject to the negotiated contract. What we have heard is that Verizon is refusing to place orders for services like Voice Over IP, Private IP and Access — services that businesses need and take time to implement. Ironically, even with VRD, Verizon will NOT be required to deliver these services any faster under the new VRD than under your existing contract.

As business pressure builds, what are you to do?

First, explain the risks to your colleagues. If the company places the order under Verizon’s new VRD contracting model, the VRD attachment terms apply in lieu of the key terms the business spent time negotiating in its Verizon master agreement and

service-specific attachments. Here are some key examples:

- Charges for the services
- What contributes to purchase requirements/minimum commitments
- Time to dispute charges
- Time within which Verizon may bill charges
- The size of early termination charges and when they apply
- SLAs

If this doesn’t convince them to slow down nothing will.

Second, just say “No.” If Verizon wants the business, it should be willing to live with the terms it negotiated. And create leverage — consider moving or threatening to move existing business (e.g., Internet access) to another provider that doesn’t dictate new and terrible terms mid-deal. This is particularly effective if you have existing contracts with other telecom providers. Tell Verizon they may not be included in future requests for quotes. Escalate within Verizon. The individuals refusing to place the orders may have limited authority to address your concerns, but higher ups can do so. Verizon will not want to indefinitely postpone new business and should be reminded of the irony if its “rapid delivery” approach is delaying new revenue.

We commend Verizon for trying to tackle one of its customers’ serious issues, but this time the carrier got it wrong. Verizon’s Rapid Delivery makes great promises and claims to solve customer frustrations, but once again Ben Franklin’s age old aphorism applies. “Take time for all things; great haste makes great waste”. ☎

¹ Verizon claims its new VRD platform yielded improvements of 70% in quote to delivery times, http://i.crn.com/custom/Verizon_PPG_Gatefold_Advert.pdf, but the improvements were only claimed for channel partners. So, before accepting Verizon’s claims, have your account representative point out anything in the VRD contract that requires Verizon to improve delivery times and watch your account rep squirm.

As of July 16, 2014, Verizon had “optimized”:

- Access (including Ethernet access and network services local access)
- Business Connection
- IP Contact Center Services
- Customer Premises Equipment and Related Services
- Ethernet Switched E-LAN
- Ethernet Switched E-Line
- Internet Broadband
- Internet Dedicated Services
- IP Business Bundle
- Location Data Service
- Mobile Workforce Manager
- Managed Call Recording
- Managed Global Network
- Managed WAN
- Managed LAN
- Managed WAN Optimization
- Managed Wireless LAN
- Network Discover
- Secure Gateway
- Wi-Fi for Business
- Application Assurance
- Mobile Merchant
- WAN Analysis Reporting
- Wavelength Service Solution
- Private IP Service (Layer 3)
- DoS Defense
- Managed Certificate Services
- Managed Security Services (Cloud and Premises)
- Security Management Program
- Security SaaS
- SSL OnDemand/Corporate ID
- Voice Over IP

Level 3’s Turnaround and the Search for the Next First-tier Provider

David Rohde

Nobody is getting out of this decade alive without a big transformation in their landline network, which is still the guts of mission-critical enterprise communications but won’t have anything to do with TDM technologies like POTS and T-1’s once we’re done.

That’s scary. But it’s also true that technology changes on this scale are a huge opportunity for end users — to freshly bid out their networks with dramatically less advantage for their incumbent national providers because those guys (basically AT&T and Verizon) won’t be able to get away with proposing the same network with a nominal price reduction.

The real scary decision here is who to put on your bid list for an RFP of such significant scope and stakes. You don’t just need “second tier” carriers to tactically play off AT&T and Verizon, sincerely or otherwise. You need qualified, potential first-tier-capable providers to whom you can credibly threaten to award the bulk of the business. And by “credibly threaten” we mean that you may actually do so.

That’s why the turnaround of Level 3 and its emergence as a legitimate candidate for large-scale, large-enterprise business is one of the most fascinating themes to emerge from 2014 and sits high on everyone’s watch list in 2015.

The very term “2014” used to give Level 3 nightmares, as it was once the year in which its heavy debt burden was going to crest — with over \$3 billion in maturities — and send it running to bankruptcy court.

The key to avoiding that was brilliant work by two executives. Level 3 CFO Sunit Patel set out in 2011 and 2012 to convince bondholders to dramatically lengthen their maturities (in fact, virtually all of Level 3’s debt now matures at the end of this decade or well beyond). Newly installed CEO Jeff Storey, an industry veteran who goes back to the “original” WilTel and its early and popular enterprise frame relay service, upended Level 3’s focus from wholesale (at one time two-thirds of its revenues) to enterprise (now almost three-quarters of revenues).

But more was needed, and in June came a masterstroke: Level 3 announced it was buying TW Telecom. Never saw that on the front page of the Wall Street Journal? That’s the point. TWT ate “boring” for breakfast and never did anything but build out metro networks, to the point where it was the non-incumbent national champion of on-net building entrances with over 20,000.

Many people (fairly) questioned both whether Level 3 could afford to buy TWT, and whether Level 3 would handle the integration well. But TWT’s own solid financials meant they were bringing significantly more profitable revenue than additional debt to Level 3’s books.

And Storey, piloting a company notorious for past sloppy acquisitions, laid out a careful integration plan that refuses to settle for redundancies and channel conflict. In fact, he noted in an early 2015 conference call that one reason TWT’s products and network management systems are so clean is precisely because it didn’t do acquisitions.

For TWT’s part, why did it even sell? Because the price was attractive, obviously — but also because TWT on its own wasn’t reaching large enterprises the way Level 3 was, particularly with its

SIP trunking service. That said, SIP trunking isn’t enough anymore. To play in AT&T and Verizon’s league the newly combined Level 3 has to show competence, scalability, product diversity, and sensitivity to individual enterprise pain points. And that’s in addition to the table stakes of strong financial bids in enterprise RFPs and a need to step up to best practices in business and legal terms — often the downfall of pretenders to first-tier enterprise carrier status.

In blog posts and at conferences during the summer and fall I laid out some specific landmarks to watch out for as this carrier emerges from its two main parts. By far the most significant was Level 3’s network investment plans. On its own, TWT was among the most liberal carriers in terms of “capital expenditures” as a percentage of revenue, but Level 3 was one of the stingiest at only 12% of revenues. To some extent Level 3’s wholesale business held back its capex — carrier-to-carrier margins really are terrible, and the wholesale business doesn’t call for nearly the granularity of building entrances that enterprises — especially verticals like retailers, hospitality, personal insurance carriers and others — demand. Still ...

I said that Level 3 would have to raise its capex-to-revenues ratio to 15% to show that it was serious about continuing TWT’s philosophy of building out to customers sold or won in competitive procurements to fulfill a strong preference for proposing Ethernet access rather than T1/T3. Lo and behold, early this year Storey said that it was doing exactly that — projecting precisely a 15% capex-to-revenues ratio for the calendar year and an intention to step up building additions to 3,000 a year.

A number of circumstances are working in Level 3’s favor as it attempts this transformation. It doesn’t have wireless network investment and the absolutely massive cost of new wireless spectrum to contend with — now a major issue for AT&T as it attempts to justify a \$3 billion cut in corporate network spending in 2015. It long ago recognized that the game is very much in the middle metro markets where enterprises won’t accept massive holes in responding to their demand sets — TWT brought 85 markets, not just the cliché “NFL Cities,” to the table.

And among the natural other first-tier contenders — the three companies other than AT&T and Verizon with ILEC territories of the sort neither Level 3 nor TWT has ever had — only CenturyLink rivals Level 3 as a national contender. Frontier (new slogan — “Everywhere you don’t want to be”) and Windstream have pursued other priorities.

Continued on page 24

WHAT IT MEANS FOR 2015

The emergence of Level 3 as a credible competitor for AT&T and Verizon in the enterprise services market is one of the good news stories of 2013 and 2014. At the very least it makes up for the disappearance of Sprint. It is distinct from, though related to, the growing deployment of SIP trunks — a market in which Level 3 is a significant player. Together, these developments promise enhanced competition and suggest that the bell is tolling for the ILECs. To most enterprise users that can’t happen soon enough.

Gigantic Spectrum Cost is Making Wireless Commoditization a Pipe(-less) Dream

David Rohde

"From sea to shining sea." That's why wireless procurement in the U.S. is so painful. Maybe Mr. Jefferson shouldn't have sent Lewis & Clark or bothered with the Louisiana Purchase.

Ideally the U.S. national policy of preserving four national wireless carriers – exemplified by the repeated refusal of the FCC to let any of the other carriers buy T-Mobile – should result in something like the commoditization of minutes and data as we saw years ago with “long distance.”

But events in 2014 demonstrated why even sophisticated users do not feel they have four co-equal choices and why carriers in general – but Verizon in particular – feel they only have to bargain so far to gain the business. That’s even in the face of what the media brands a “price war” in the consumer market, something of a half-truth based on advertising come-ons and something that actually threatens to ruin, not enrich, the most “aggressive” carriers.

The basic problem here is that the United States is too big. Strike that – the U.S. is fine as it is. It's just very, very large geographically. It's easy to forget, especially living or doing business anywhere in the East Coast's I-95 corridor, that the U.S. is probably the least densely populated of all “First World” or fully industrialized nations other than Canada and Australia. Covering the country requires an absolutely enormous amount of spectrum.

Compounding the issue is how much wireless changed culturally in the first half of this decade. I'm not talking primarily about how reliant people are on wireless data and how much power users insist on video and other apps working just fine on their devices. I'm actually talking about "coverage forgiveness," something that has virtually evaporated as a social concept.

In 2010 people could travel and discover that their cell phones don't work in another part of the country, and shrug. In 2015? Forget about it. Mobile phones are supposed to get you a signal just like landline phones get you dial tone. End of story.

Following an extraordinary \$45 billion bid for new spectrum in the AWS-3 auction that straddled 2014 and 2015, and with a critical auction of low-band, high-propagation spectrum commandeered by the government from UHF television stations scheduled for early 2016, the question often comes down to who among the carriers can afford to stay ahead of the game.

Or stated simply, why does Verizon

- Thumbing their nose at the enterprise wireline market without realizing what that cost them in terms of “share of mind” in large businesses as a strategic vendor in general.

That helped Verizon pair itself with AT&T as top dogs rather than the old pairing of “MCI” with Sprint as scrappy contenders to the AT&T throne. In fact, arguably many Americans now consider Verizon the top telecommunications company by brand value, not AT&T – but with everyone else way, way behind.

So let's look ahead. In early 2016 some of the 600 MHz spectrum that has been used as UHF spectrum by TV stations is due to be auctioned off to wireless carriers in what's known as an "incentive auction." This auction will have mind-numbingly complex rules and phases that have been proposed and debated.

T-Mobile says and does – particularly its fascinating CEO, John Legere.

On the one hand, T-Mobile clearly believes that it was outgunned in the AWS-3 auction by Dish Network by Dish setting up multiple bidding entities that were colluding to drive up prices beyond T-Mobile's reach. On the other hand, T-Mobile often gives the impression that it wants to merge with Dish in order to give T-Mobile a further reach into the consumer market (while giving Dish what it wants – wireless data to offer consumers and businesses and to get away from the “dying” market for “television” programs that millennials don’t watch).

It's fascinating to observe Legere's own Jekyll & Hyde routine here. Just call up his Twitter feed to see his multiple and contradictory comments about Dish. But Dish has a trump card – the circumstantial fact that it is NOT currently a national wireless carrier. The two previous bidders for T-Mobile – AT&T in 2011 and Sprint (for all intents and purposes) in 2014 – were foiled by federal regulators because their takeover bids would have reduced the number of U.S. national carriers from four to three. Dish doesn't have that “antitrust” problem and can easily argue that merging with T-Mobile makes the consumer wireless market more competitive than ever. However, your problem (and ours) with Dish and T-Mobile merging would be that it would disorient T-Mobile from another important priority: focusing on the enterprise wireless market as much as Verizon and AT&T do.

In short, every competitor to Verizon has a spectrum and network-investment-related challenge to face in order to convince you to buy their services when they bid lower. AT&T is essentially a peer but it has had to go into the bond market to fund its \$18 billion in AWS-3 spectrum purchases, suffering a credit rating downgrade in the process – meaning it may be gun-shy when the 600 MHz auction comes around. Sprint is losing money big-time and its brand image has taken an enormous hit, all while continuing to score low in network quality tests despite recent upgrades. T-Mobile's aggressiveness has been welcome and fun to watch but it still has to fill in small-town coverage across the enormous North American continent in the 600 MHz auction, and it basically has to decide what it wants to be when it grows up – a consumer value play or a serious enterprise contender. The wireless game in 2015 and into 2016 is not only critical for enterprise telecommunications managers, but quite simply one of the most fascinating stories in American business, period. 📞

The basic tension in these decisions about the rules is whether AT&T, Verizon, Sprint and T-Mobile will all have equal bidding rights to the spectrum, or whether Sprint and T-Mobile will have a leg up.

During 2015, much of the ongoing debate over the rules for the 2016 incentive auction will be based on claims and counterclaims by the four national wireless carriers over who “gamed” the AWS-3 auction – including Dish Network, which desperately wants to diversify beyond delivery of “television” programming and get into wireless data. In particular, watch what

often get away with charging more for “network quality”?

Part of it is because a little bit of luck came Verizon's way, especially in the form of horrible strategic errors by Sprint that dug a hole they may never really climb out of. Notably two:

- Delaying LTE deployment because of a pipe dream around WiMax fixed wireless access built on some especially dodgy partnership deals.



How to Avoid eRate Rule Violations

Steve Rosen

The FCC is investing millions of dollars to remove eRate waste, fraud, and abuse. Here’s how to make sure you’re not caught in its net.

This summer, the Federal Communications Commission (FCC) created a Universal Service Fund (USF) Strike Force, which is tasked with combating waste, fraud, and abuse in various USF programs, including the eRate.

This newly created Strike Force, which is part of the FCC’s Enforcement Bureau, will almost certainly expend considerable resources ensuring that the procurement practices of schools receiving eRate funding comply with FCC rules.

To avoid a future encounter with the Strike Force, schools should re-evaluate their internal compliance programs—and here’s how.

For schools, one of the biggest challenges in the eRate process has been complying simultaneously with state and local procurement rules and with the separate, occasionally inconsistent requirements of the eRate.

While the eRate requires adherence to state and local procurement laws, the program also requires applicants to comply with a number of rules that are unique to the program, or face a denial of funding. Because these eRate rules are not always ingrained into a school district’s procurement practices, they are sometimes ignored or misinterpreted—with disastrous results.

There are as many opportunities to violate eRate rules as there are rules, but the most common rule violations are:

- Failing to create a written bid evaluation matrix.
- Failing to make the price of eligible services the most heavily weighted factor in the bid evaluation matrix.
- Failing to wait 28 days after posting a Form 470 and before awarding any contract to a service provider.
- Submitting a Form 471 to the Universal Service Administrative Company (USAC) before entering into a contract with the service provider. While USAC defers to state law as to when a contract has been formed, it does insist that this formation occur prior to a Form 471 being submitted.
- Choosing a service provider for the installation of internal connections without obtaining a price for all of the parts and labor associated with this installation. Schools

are not permitted to hire an entity as a “general contractor” and then have that entity procure the parts and labor necessary for the installation.

The FCC prohibits such procurement practices because they put an entity other than the school district in charge of the project. Moreover, the “general contractor” approach does not allow schools to compare each bidder’s price of eligible services, which is central to the program’s procurement rules.

In the face of all of these eRate rules, how are schools to stay on the right side of the law? The best way to minimize your risk of compliance problems is to take some or all of the following steps:

- Empower a core group of employees to oversee compliance with the eRate. Because the program’s rules are complex and often counterintuitive, it is rarely beneficial to allow untrained (or, perhaps worse, partially trained) employees to interact with the program. A small group of employees who develop subject matter expertise regarding eRate compliance can literally save you millions of dollars.
- Keep that core group current on changes to the program’s requirements. Either prepare periodic in-house training sessions, or take advantage of USAC’s online materials—or do both.
- Prepare a written eRate compliance manual that all employees involved in the program can reference, and update this resource regularly.
- Appoint an in-house eRate expert who can troubleshoot any eRate issues and serve as a resource for other employees with questions about the program.

The eRate is an excellent source of federal support for internet-related services. However, the program is filled with a number of complex—and often counterintuitive—rules that, if not followed, can lead to a visit from the Strike Force, a loss of funding, and all the bad things that follow. With that in mind, applicants should ensure they have employees who are well versed in the program’s rules. 📞

Continued from page 6

customers abandon traditional (contributing) telecommunications services for Internet-based (non-contributing) telecommunications services. If the FCC added broadband Internet access revenue to the base, the factor would drop below 5%. But Congress has complained that the FCC would be “taxing the Internet” so the FCC isn’t going there. As a result, the disproportionate share of USF paid by enterprise customers will continue to grow.

There’s another rulemaking at the FCC that is examining the current USF contribution methodology and could be used to fix the USF problem with new rules. Perhaps when the political firestorm around net neutrality dies down, the FCC will be able to turn its attention to the broader USF contribution proceeding and amend its rules to make them more equitable. Savvy enterprise customers should participate in that docket. 📞

Continued from page 20

Frankly, one of Level 3’s biggest challenges is to fend off the typical Wall Street framework where analysts press companies for more merger “synergies” – the thing that looks lovely on stock analysts’ spreadsheets but in real life are experienced by customers as the endless screw-ups that bedevil the typical telecom merger. I’m here to report to you that in the earnings calls for the third and fourth quarters of 2014, Storey gently but firmly told analysts to buzz off about the issue of savings beyond what Level 3 had originally promised in its merger filing. Good for him.

Much remains to be seen — only bids on IP transformation RFPs will really tell the tale. But it’s a helluva turnaround to be talking about Level 3 this way. Enterprise users are eager to see the story continue to unfold. 📞

Please note that the articles in this 2014 Year in Review have in some cases been edited from the originals.

Negotiating Private Cloud Transactions: Building the Foundations of Best in Class Deals: Part 1 – Introduction and Pricing and **Part 2 – Performance, Control, and Risk Allocation** (Ben Fox and Marc Lindsey, 4/7/2014) were originally published by No Jitter (nojitter.com)

SIP Trunking is Driving Users’ Ability to Consolidate Vendors in Europe (Ben Fox, 2/19/2014) was posted in TC2’s David Rohde on Telecom blog (techcaliber.com/blog).

The following articles are concepts based upon compilations of original posts (techcaliber.com/blog):

- **Level 3’s Turnaround and the Search for the Next First-tier Provider** (David Rohde)
- **Gigantic Spectrum Cost is Making Wireless Commoditization a Pipe-(less) Dream** (David Rohde)

The following articles were originally published by CCMI (ccmi.com):

- **Why Ethernet Access is a Critical Part of Enterprise Wireline Networks** (Hank Levine and Jack Deal, 6/2/2014)
- **Verizon Rapid Delivery – Not SO Fast** (Laura McDonald and Deb Boehling, 9/1/2014)

The following articles were originally published in Webtorials (webtorials.com):

- **Enterprise Sourcing of SIP Trunking Services** (Hank Levine and Jack Deal, 2/12/2014)
- **Three Things You Really, REALLY Need to Know Before Tackling SIP Trunking** (Hank Levine, 5/13/2014)
- **AT&T Issues Withdrawal of Service Matrix – Part 1** (Deb Boehling and Janis Stephens, 12/16/2014)
- **AT&T Issues Withdrawal of Service Matrix – Part 2** (Deb Boehling and Janis Stephens, 1/8/2015)
- **What the FCC’s “Open Internet” Ruling REALLY Means for the Enterprise** (Colleen Boothby, 2/25/2015)

LB3 and TC2 assist enterprise customers with network service and IT procurements, benchmarking, compliance management, regulatory issues and disputes. Learn more about industry developments and the challenges facing enterprise users.

Levine, Blaszak, Block & Boothby, LLP ("LB3") specializes in telecommunications and technology law, with particular emphasis on the representation of large users, including almost one-half of the Fortune 100. LB3 has extensive experience in negotiating custom network service agreements, network outsourcings, and related transactions on behalf of enterprise customers. Together, LB3's 14 partners have assisted large users in connection with over 1,000 network services agreements and related transactions and have more than 225 years of experience in the field. Beyond transactions, LB3 is the leading representative of large end users and IT companies before the FCC and other regulators, and is the first choice of large end users whose relationships with their communications providers have broken down. LB3 also advises clients in connection with software licenses and telecom-related acquisitions.

TechCaliber Consulting, LLC ("TC2") helps major businesses get the most out of their voice services, data networks and managed services at the lowest possible cost. TC2 cuts through carrier efforts to obscure prices by establishing target prices before negotiations begin using a wealth of hands-on experience in the market for corporate voice and data services and managed network operations. TC2 works with clients to support and/or lead the proposal process and negotiations, enhanced by its affiliation with LB3. Most importantly, TC2 helps clients structure their obligations to maximize each client's permanent, ongoing leverage as real prices decline, new technologies emerge, and telecom industry upheaval reshapes both domestic and international procurements. TC2 works primarily for large multi-national companies, supporting over 100 deals each year that range in value from a few million dollars to over \$1 billion.