

In the Spotlight: Kamran Sistanizadeh, Yipes

By Ed Gubbins

Aug 22, 2006 3:51 PM



Before teaming with RAD Data Communications to [extend](#) Ethernet services across DS-3s and OC-3s, Yipes made news by offering [microsecond Ethernet service level agreements](#) (SLAs) using HawkEye, the performance measurement system the company developed in-house. Yipes co-founder and chief technology officer Kamran Sistanizadeh recently spoke with *Telephony's* Ed Gubbins about managing carrier Ethernet services and SLAs and where the industry's habits in that regard are headed.

On the new SLAs: These SLAs provide a launching pad for some of the value proposition we offer. Customers are looking at these SLAs for mission-critical applications. They now know on a piece of paper the performance is legally guaranteed. Especially in time-sensitive applications such as trading, we've certainly been quite successful penetrating financial verticals. Having layer 2 performance SLAs is very important in the eyes of brokers and traders. Anyone who has applications that are time-sensitive can now use our layer 2 network.

On [Hawkeye](#): The HawkEye system is not, per se, a property of the service itself. It's a measurement tool. It captures information from the devices we have placed on performance and provides a consolidated view to end users and to the operational organization in the company. We've not changed the boxes or the equipment. We created HawkEye to capture the SLA information and present it to the end user. Even if we didn't have HawkEye, the salient properties of the network would still not change. It's not the HawkEye system that removes jitter from the network. It provides the end user with the properties of the network. The [equipment] provides generic raw data out of these devices. Every box has [a management information base] that provides certain information about the behavior of that box in its own location. When you build a network, you have to create a relationship between these devices and map them into the services you offer. There has to be a lot of customization beyond buying and operating the boxes. That customization is specific to every service provider. There's no such thing as an out-of-the-box equivalent of the HawkEye system. It has to be fully customized.

On [virtual private LAN services]: VPLS is kind of a pivotal driver in our Ethernet strategy, creating a completely managed, fully meshed, Layer 2 backbone network. In all our major markets, we have placed VPLS devices and created a fully meshed VPLS network across our backbone. It's an area that really brings out the managed Ethernet service value proposition on a global basis. It's a tool that allows any customer not just within a metro but a global area to be part of a switched LAN cloud. Any customer can enter its port, leave its port and create connectivity without incurring loss of operations or capital cost.

On capacity engineering: Those SLAs are dependent on the way we designed the network. We capacity-engineered the network. We provide customization on the software that runs the network; we manipulate the code that runs the devices. There are a lot of IP tricks you can play with devices, and the code allows you to differentiate yourself in the market. From a capacity perspective, ensuring a sufficient level of redundancy and resiliency, you have to make sure you don't oversubscribe, that you're not losing packets. When a customer buys a 100-Mb/s service, you have to make sure it's available at every single entry and exit of the route. Everything has to be engineered to provide the equivalent of a private line without losing a packet. These are all bits and pieces of the design of the network itself as well as the definition of the service. Anyone can go buy boxes from Cisco [Systems] or Extreme [Networks]. The trick is, when you get these boxes, what you do from an engineering design perspective to connect these boxes and run a service on it. That is the differentiation that a service provider brings. After six or seven years of doing this, we have the internal know-how to do this.

How you manage the capacity and flow of traffic and how you manage and utilize links in an optimal manner to reduce internal costs is different than saying capacity or the network is oversubscribed. When someone hears that a network is oversubscribed, that means when he sends his own packets, they're probably not going to reach the destination. We don't do that. We capacity-engineer the network such that we don't drop the packet when the network is being utilized by a certain application. If a customer says their application is only going to run for 6 hours between midnight and 6:00 in the morning, that application is going to run during that 6 hours, but during the remaining 6 hours, that link is empty. So maybe I can use it for some other customer. During that 6 hours, I have to make sure that network is not oversubscribed. There's a subtle difference here between the notions of over-subscription and capacity engineering. If properly capacity engineered, from the customer's perspective, it's never oversubscribed. If the network is oversubscribed, the customer doesn't know.

On T-1s vs. Ethernet: T-1s have been around for almost 40 years. It's a very mature, well-penetrated technology. Anywhere you go across the globe, if you say, 'I want to order a T-1 line,' or an E-1 line in Europe, everybody understands what you're talking about. There's a very well-established standard and rules of engagement for ordering T-1s. In the case of Ethernet, from an interoperability perspective across various carriers and the ease of connecting different networks together, there's still work that needs to be done. Just because we have boxes and

vendors providing engineering requirements doesn't necessarily mean that service provisioning is as easy as just buying a box. Work still needs to be done in service provider organizations as well as across service providers to make sure network-to-network interfaces and service provisioning flow between service provider A and B. I'm assuming it's not going to take 30 years, as it did with DS-1s and DS-3s. It's going to take probably another couple years for interoperability issues to be sorted out. That's not a technology issue. That's about business and operational processes. It's about the handoff between service providers. It's a matter of time before ease of use becomes clear.

On SLA differentiation: We try to be a little bit ahead of that. It's a matter of understanding what the industry as a whole is gravitating toward. If there are changes we have to make in our network to meet those requirements, we will, but usually in a standards body, most of the items that need to be resolved get into a common denominator, so a lot of people agree on it. In some cases, it loses that level of differentiation because there are a lot of parties involved. To the extent we feel it's appropriate to offer exactly what the standard says, we will, but we like to always be ahead of the curve and provide a superset of what those standards say rather than a subset. If I'm providing QoS or SLAs that are better than some other service providers' and the customer is willing to buy it, I should not wait for the standard. We have to listen to our customers. If the standard is there, that's fine. If it's not, we have to customize and move on. We follow the standards because the equipment providers follow the same standards. But at the end of the day, the intellectual property is not mandated by standards but rather by the capital and the intellectual organization itself. That's how we want to be.

On suppliers: We have multiple layers of technology within the network. We still use Extreme [Networks] equipment in our metros. Juniper [Networks] equipment has VPLS components used at the edge of the metro where we want to connect metros on a global basis. We have a layer 2 Ethernet environment in metro areas equipped by Extreme. Then we extend this to another metro using a Juniper backbone.

On Extreme's Black Diamond 12K: It's certainly on our roadmap for complete evaluation. We're now evaluating it and will be field trialing it very soon. Larger companies like Verizon [Communications] and AT&T have legacy networks and are going to take a different path of evolution than a company like Yipes who started with Ethernet and doesn't have legacy services. When someone says is Black Diamond big enough for you, do you need to push MPLS inside it or inside the network, it depends on how the network was designed originally and what is the purpose of that evolution. I don't believe I have a scalability problem based on the capacity these devices offer. I'm a much smaller company than Verizon or AT&T and don't have those scalability issues. Am I going to be able to support millions of customers in a metro? No, because I'm not AT&T or Verizon.