OPASTCO
ROUNDtable
BIG IDEAS FOR SMALL

SPECIAL
BUYER's GUIDE SECTION

HATS OFF
to Speeding Up Broadband Deployment and Reaching New Public Relations Audiences

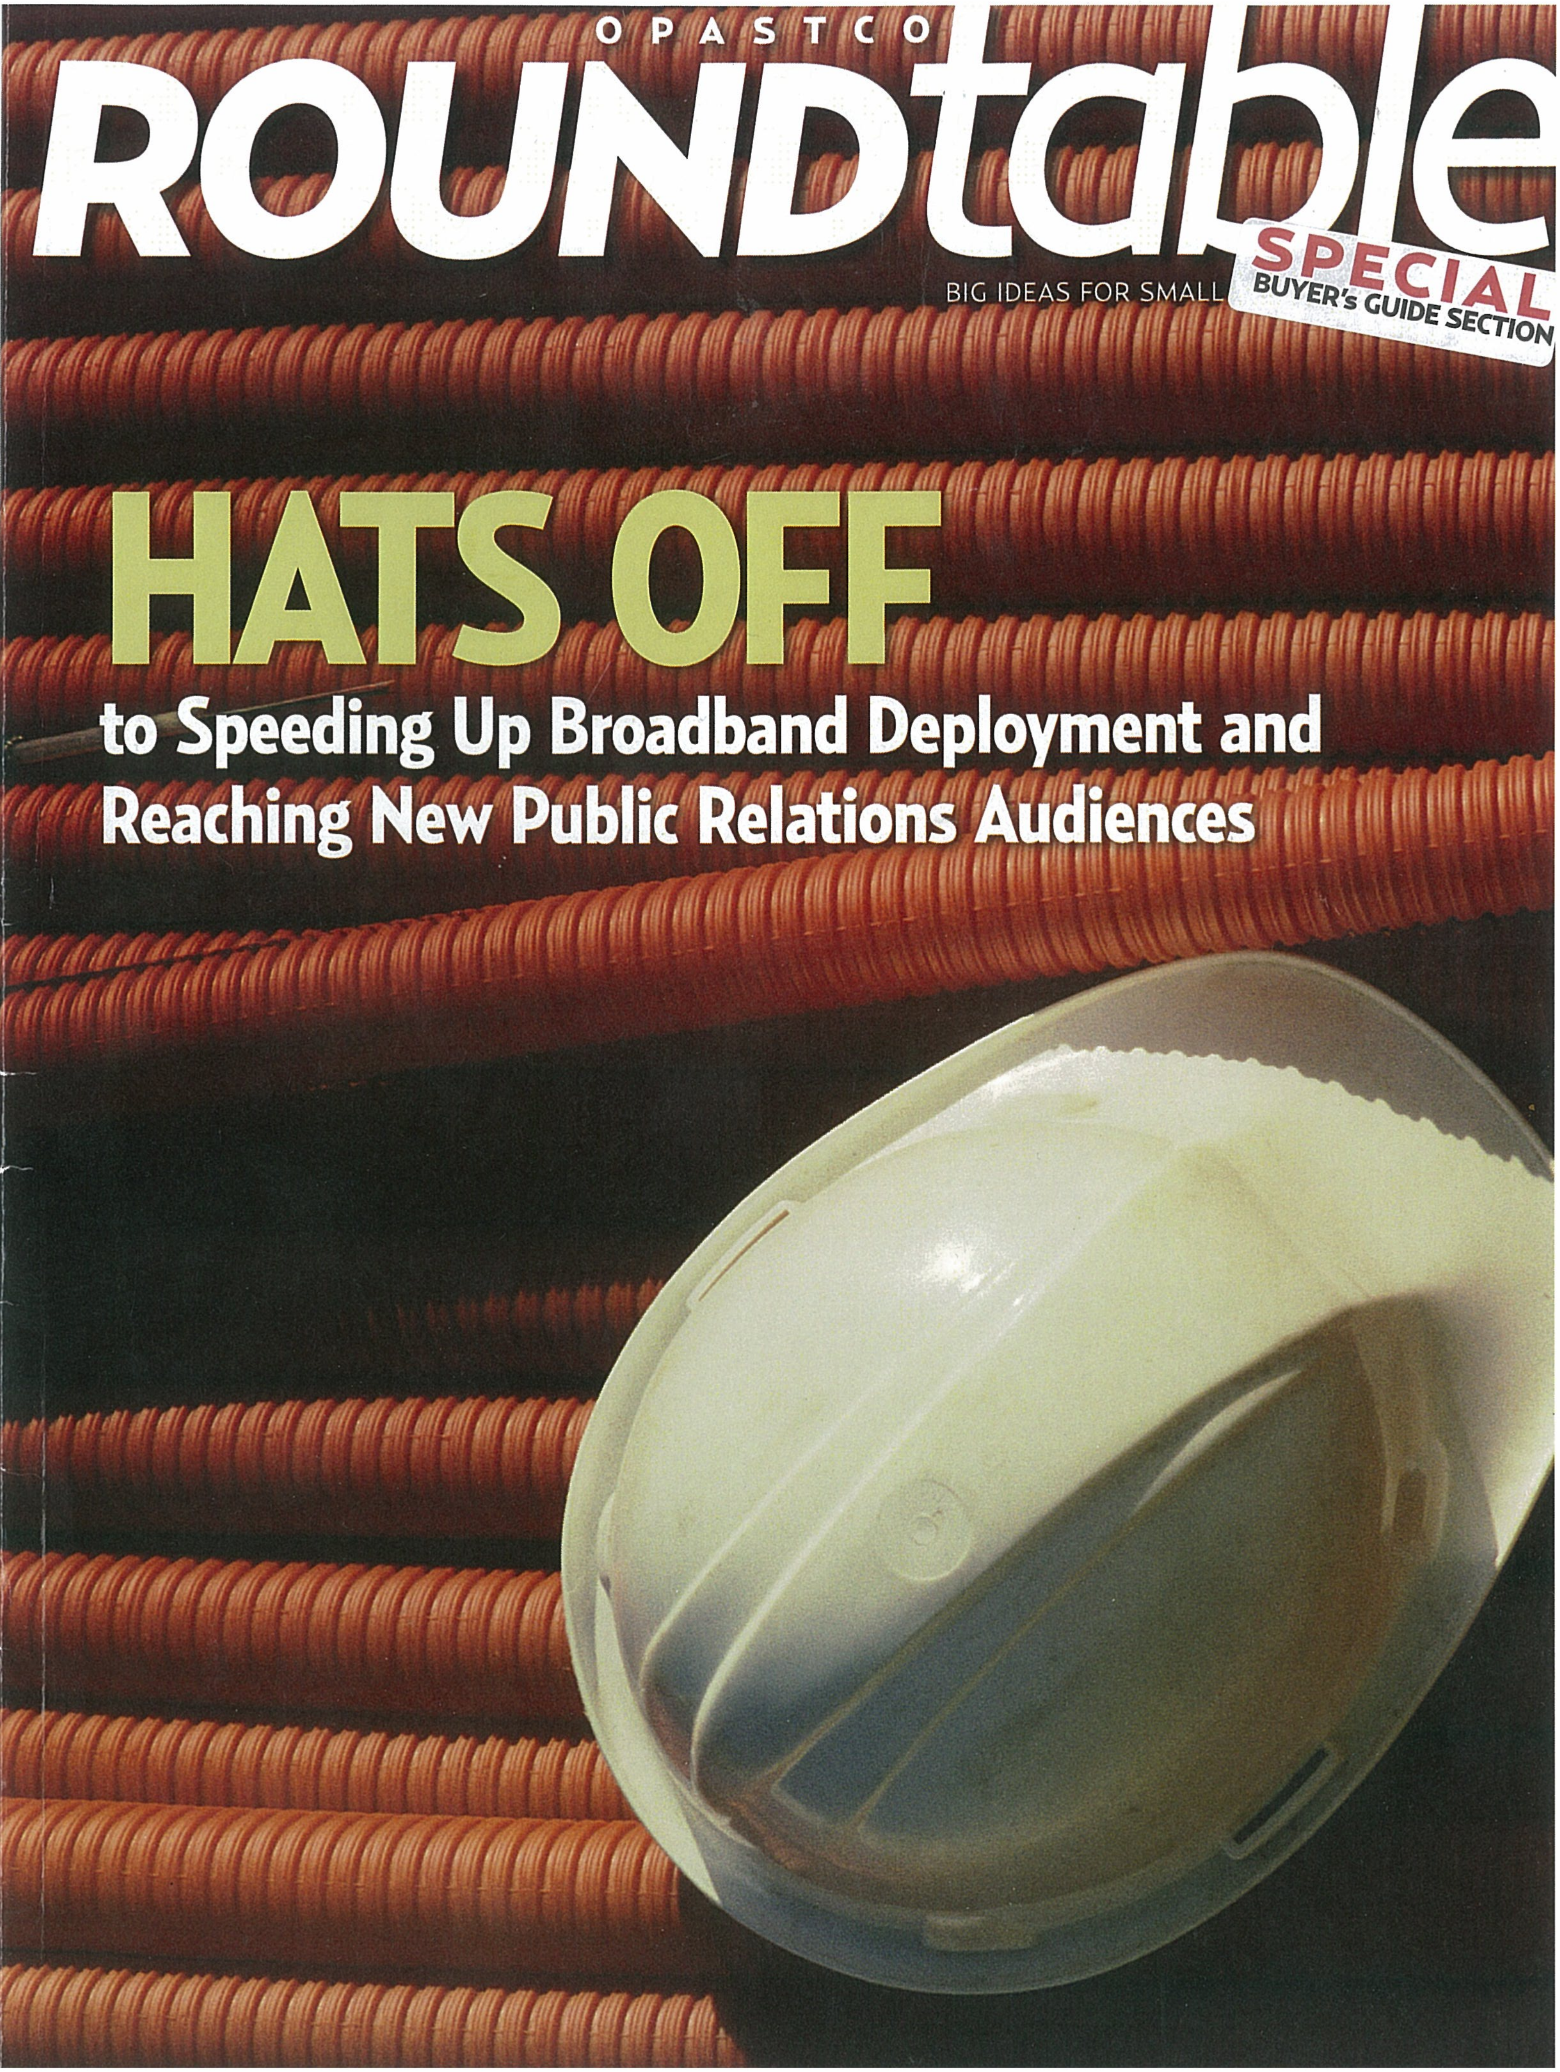

Understanding The Separation Of Access & Transport In Packet Based Networks

Access and transport via Ethernet has several advantages worth considering for your broadband network

MIKE SKUBISZ

As the proliferation of bandwidth-hungry applications continues, network operators are quickly re-engineering and deploying new network access and transport platforms. These new platforms are, for the most part, built around IP over Ethernet technologies. One of the key questions that arise with the use of Ethernet for both access, as well as transport, is where does the access network end and the transport network begin?

Just a few years ago there was a clear, physical separation between access network equipment and high-speed transport equipment. Even with the most popularly deployed access equipment, the technologies used on the backplane were different than those used for transport and access.

A backplane is the connector board usually found in the back of an electronics frame, where cables terminate and modules plug into. The backplane is the determining factor on how fast the modules and plug-in can operate.

Today, the physical separation that once existed between access network equipment and high-speed transport equipment is no longer so clear.

Standardizing on a Single Technology

Many of today's modern, packet-based telecom platforms are utilizing Ethernet as the protocol of choice for transport, as well as access, and even as a backplane technology.

This rapid shift to a standardized access protocol is hardly revelatory. We already bring Ethernet into our businesses and homes. It's not that far a leap as a preferred transmission protocol as well.

There are several benefits that accrue to the telco operator in standardizing on a single protocol for access and transport. Because it's used so widely, Ethernet becomes the best price/performance technology available. Secondly, almost every piece of content that enters a telco network originated from an Ethernet-sourced device and is also delivered to the customer, ultimately, in packet form. Even if you're using ADSL, for example, the modem will convert the signal into Ethernet. In sum, using Ethernet to satisfy both access and transport requirements for a telco provider is less complicated, less cost prohibitive, and presents fewer barriers to interoperability among network-wide devices. It also is technically competitive as a protocol with what it typically replaces for transport applications, e.g. Synchronous Optical Network (SONET).

Separation Then, Integration Now

Paralleling the emergence of Ethernet and packet-based networks in the telco community is the integration of access and transport in the same Multi-Service Access Platform (MSAP) chassis.

This meshed approach to transport is not only more efficient than traditional technologies like SONET, where access and transport were physically separate boxes and required separate support and management resources, but it is also far more scalable, reliable and cost-effective.

A CAVEAT: Just because access and transport now use a common packet-based technology, realizing the most benefit from a unified solution protocol requires you, as an operator, to operationally separate the transport pieces of your network from those supporting the access parts of the network. It's a subtle but profound difference. For example, chances are you're not running certain features on the access side that you're running on the transport side, and there are certain features on the transport side that are probably not running on the access side.

Practically speaking, if one of your access modules fails, a handful of your customers might lose connectivity; if you have a transport failure, you've may have just brought down your entire network.

Separate But Equal

With access and transport connections built into the MSAP chassis and an orientation to regard them as separate operational areas of the network—primarily as the result of sharing common technologies, albeit treated, maintained and supported differently—the next step is to decide which backplane "approach" you favor.

The previous generation of telecom devices typically used ATM or TDM as its backplane protocol. These first generation backplanes were principally designed for carrying copper access technologies (DSL and POTS). At the time, that particular backplane was sufficient for the type and amount of traffic it carried, in this case mostly voice. However, with the primary class of traffic traveling across today's backplanes being high-speed Internet and Internet protocol television (IPTV), there has literally been a tectonic shift away from ATM or TDM to packet-based technology.

There are two different popular approaches to building out an Ethernet backplane to reduce transport failure and optimize access card availability and performance. The first approach physically, and operationally separates access from transport, while a second integrates both transport and access onto individual cards that are physically as well as logically serial stringing one to the next.

In the instance of the former chassis configuration, its backplane uses packet-based technology and contains a pair of dedicated line (or transport) cards, designed to take in the fiber optic transport connection from the outside world and to provide transport protection in the event of a fiber optic cut or equipment failure. The remaining slots are populated by access cards that connect to every other slot in the system. However, they are not in the physical path between the transport cards. The distinction is important. In this non-blocking backplane architecture, the access technology is separated from the transport technology. As a result, they operate independently such that if you were to physically remove one of

"Practically speaking, if one of your access modules fails, a handful of your customers might lose connectivity; if you have a transport failure, you've may have just brought down your entire network."

> **"For chassis where access and transport are integrated onto individual cards, it becomes unavoidable for network operators to double or even triple the number of chassis on their ring in order to realize full redundancy and maximum performance."**

these access cards, the transport remains unaffected. Conversely, how well or at what speed the transport cards work are transparent to the access cards residing in the chassis. In other words, access cards can be pulled out (whether intentionally or otherwise) and only those users assigned to that card are affected. The network, again, remains intact. In fact, being physically and operationally separate, any access card can be any speed, regardless of the speed of the transport card itself. This backplane configuration lowers costs, while maximizing backplane performance.

But what happens when every card in the chassis is both an access and a transport card?

In this case if an access card fails, the entire network is affected because both logically as well as physically, it really is all one network. By extending the transport across the entire backplane every line card becomes part of the transport process. This approach becomes increasingly problematic as any time a line card is pulled in this configuration, whether for upgrade purposes or to move users, the entire network is affected and performance is compromised. As you can see, the methodology of serially stringing the transport across your backplane essentially exposes the backplane and transport to far greater potential failure.

Put it another way, in our first chassis example if you as an operator want to run your transport at 10 GB, feel free. The cards themselves can run at any speed you want. Although the systems are separate and there are inherent dependencies between them, they are still operationally unique from a transport and access perspective. Indeed, if your access cards need to run at 10Mbps or 100Mbps or even 10Gbps, they can—it has nothing at all to do with how fast your transport is running.

Conversely, in our second chassis example, where transport and access are meshed into the same card, your network becomes as slow (or as fast) as your slowest card. Additionally, if each card in the chassis has to run at 10gig each, you are absolutely carrying this cost if higher performance transport into every component of the system. In this scenario your cost, as well as potential transport failures/events, goes up exponentially.

A final thought on this particular subject. For chassis where access and transport are integrated onto individual cards, it becomes unavoidable for network operators to double or even triple the number of chassis on their ring in order to realize full redundancy and maximum performance. As a result, if you have 100 access cards with integrated transport on a ring, you essentially have 100 transport nodes. For chassis where transport and access are dedicated, you will, ultimately, have far fewer transport nodes and a much more reliable network. As a network operator, this hindsight towards building the best design possible increases redundancy and reduces Mean Time to Failure (MTTF).

Speaking of MTTF, the argument is still valid: the more parts involved in a system, the higher the probability of failure as the frequency between failures increase. In the dedicated transport chassis configuration above, this kind of system has a higher Mean Time Between Failure (MBTF), is more reliable and available and the failure of an access card doesn't trigger a transport event. Moreover, speeds equal cost. The faster something goes, the more expensive it becomes. The separation of transport and access allows you to de-couple the transport speeds from the backplane speed. What does this mean for you? Over the long term it could be mean a lot.

Over the next five to 10 years, as industry and bandwidth standardizes on 40 or even 100 gigabit transport, ask yourself the following question: do I really want to carry the cost of 100 gigabit across every line card in my system? The answer time and again should be no, I want my backplane and access line cards to run only as fast as the line cards themselves require.

A Sensible Migration Path

As telco operators migrate to packet-based networks, remember that your legacy ATM and TDM traffic is not going away. TDM is already very sensitive to delays and disruptions; therefore, having access and transport separated will likely help protect and stabilize traditional telco traffic.

As you consider operationally separating (or integrating) transport and access cards, take a moment to also consider what kind of chassis will cost effectively evolve your network to all IP without disruption to those legacy services. Evaluate solutions that are distributed and that offer multi-service access platforms built for broadband access and aggregation. Consider, too, those solutions that scale to support a wide range of access technologies, from legacy T1 and POTS to ultra-broadband, ADSL2+, VDSL2, gigabit-capable passive optical network (GPON) and active point-to-point Ethernet subscriber connections, all in a single integrated system.

In sum, look for solutions that are:

- Ethernet-based
- Operationally Separate Transport and Access
- Carrier Grade
- Cost-effective
- Flexible in choice of access method

Choosing wisely ensures a seamless evolution of service provider legacy networks to modern, efficient, carrier-grade optical Ethernet transport and access.

Mike Skubisz is chief technologist for Pannaway. Visit Pannaway at www.pannaway.com