COMMSCOPE®

Propel™

An interview with Ken Hall, CommScope Solutions Architect

What are the key drivers fueling the changes in the data center physical layer infrastructure of hyperscale and cloud-scale providers?

There's been a rapid migration to higher speed data rates over the past several years, mainly due to the need to increase capacity and responsiveness in data centers. Data center managers are looking to take advantage of faster, higher capacity switches. To do that, they need to provision more ports at higher data rates and higher optical lane counts per port. Among other things, this requires thoughtful scaling with more flexible deployment options. Simplifying and supporting the required design, installation, operations and migration paths means infrastructure and network teams must collaborate to ensure the cabling architectures align with the network configurations.

What we've found with our hyperscale and global-scale clients is that they need the right building blocks to enable a flatter network—one that delivers much better performance and redundancy. The evolution from four-lane quad designs to eight-lane octal has enabled the migration to 400G, 800G and eventually 1.6T and beyond. The 16-fiber configuration that supports octal technology is the primary building block.

This, in turn, is driving changes in network topologies. How do you distribute all that capacity most efficiently, switch to switch and switch to server? We know that the trend is to flatten the network by reducing switch layers where possible.

But each use case is different; therefore, a high degree of flexibility is necessary. That means providing the greatest range of breakout options and interfaces possible. Not only the 16-fiber building block, but 2-, 8-, 12- and 24-fiber configurations that support legacy applications. These are just some of the many changes that are helping to reshape the requirements and designs of today's hyperscale and cloud-scale data center networks. We designed the Propel™ platform to help our customers meet those challenges more easily, efficiently and gracefully.

So, how would you describe Propel?

Propel is an end-to-end, high-speed, modular fiber platform that, we believe, best addresses the current and future needs of today's larger, high-fiber count data center networks. It is built on three pillars:

Design flexibility. Propel is the first global fiber platform to incorporate 16-fiber MPO technology, which, along with 8-fiber MPO, are the building blocks for 400G, 800G, 1.6T and even 3.2T migration. Propel also supports customers' 8-, 12- and 24-fiber legacy deployments. Four different, but fully interchangeable, fiber module sizes simplify design and alignment between network applications and connector options. All fiber modules fit in a single panel design, further simplifying upgrades and changes without field modification.

Ultra-low loss performance (ULL). As networks flatten and link spans increase, ULL optical performance is critical for delivering high-speed transmissions. Propel features precision, angle-polished connector (APC) end-faces and a proprietary system for exact fiber alignment that guarantees consistent ULL signaling—multimode and singlemode—over longer distances.

Efficient Day 1/Day 2 operations. Propel's third pillar enables data center managers to deploy, upgrade and manage their fiber network faster and easier than current solutions. The common panel design offers front and rear access—making modules, adapter packs and cable assemblies easier to install and service in areas where space is tight. Polarity is standardized on Enhanced Method B, so there's no need to flip modules and assemblies, and the connectors use aligned keys—ensuring transmits go to the proper receives. The blades can be removed and re-inserted, if necessary, by a single technician. As mentioned, the modules on any blade are totally interchangeable and the blades fit into a single common panel. So, you don't have to switch out or field modify panels when you change modules. Cable management is also easier with improved routing and secure cable clips that easily snap on and off of the blade.

What is the genesis of the Propel platform? What drove its design?

Propel actually grew out of our High Speed Migration platform. Shortly after launching that a few years ago, we began working on what would become Propel. We went to our customers, sales team, and integrators and installer partners and asked them: "If you could have anything in a fiber platform, what would it be? If you're looking at the applications in the market, how would you want to support them?" Their answers echoed the same themes repeatedly—16-fiber migration, ultra-low loss optical performance, front and rear panel access, interchangeable components, etc.

Propel is an end-to-end, high-speed, modular fiber platform that, we believe, best addresses the current and future needs of today's larger data center networks.

Customer input was critical on these issues. They know our portfolios' legacy and performance, along with other products available in the market, and told us what they like and don't like. We also put those same questions to installers and integrators, the professionals who work with these systems daily. We then challenged our design engineers. "If you had a blank canvas, knowing what we know now about network design, migration and installation—as well as applications, performance and on-site time sensitivity—how would you design a solution that simplifies all possibilities effectively?" And then we weighed a variety of different scenarios. What are the critical options? How do we provide capability, flexibility, the right building blocks? The result is a big leap forward in terms of a highly flexible, future-ready fiber platform.

What are the key differentiators between Propel and the other fiber platforms?

There are a few. Certainly, the 16-fiber support is unique and ahead of the market. The platform is also application-specific so customers can easily tailor their networks to the changing demands within their data centers. That's important as network teams and infrastructure teams become more collaborative. Then there are the interchangeable modules that enable you to reconfigure channels quickly and easily, whether you need four duplex ports for an 8-fiber application, eight for 16, six for 12, or 12 for 24. You can do that inside the panel without any field modification and without switching out panels.

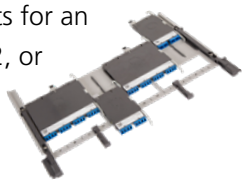

Another important difference that sometimes flies beneath the radar is that Propel customers are fully supported with CommScope's Application Assurance and SYSTIMAX® 25-year Warranty. This means your Propel links are guaranteed to support the applications for which they are designed, now and in the future.

How did you decide it was time to include 16-fiber MPO connectivity?

First, from an application perspective, you have to realize that a limiting factor is the number of I/O (input/output) ports on the front of the switch. No matter what the capacity is inside the box, you've only got space for a maximum of 32 QSFP/QSFP-DD/OSFP ports on the front of a 1RU switch. Then the question is, how do you distribute that capacity from the switch? When the industry went from duplex to four-pair quad designs, it enabled us to split out the native

Propel links are guaranteed to support the applications for which they are designed, now and in the future.

data rates—whether it was 10G, 25G, 50G or 100G—to four separate devices. With 16-fiber, we can now go from that four-way breakout to eight regardless of lane speed, whether it's 25G, 50G, or 100G lanes, and possibly 200G in the near future. That's the beauty of the octal technology.

The trend to octal allows you to flatten the network—reducing the number of layers needed deliver the same number of ports. Doubling the number of breakouts enables you to eliminate some switch layers, such as top-of-rack (ToR)—reducing the number of switches in a cabinet row from eight to two, as an example. Moreover, today's applications are being designed for 16-fiber cabling. So, deploying 400G, 800G, or 1.6T applications with that 16-fiber infrastructure distributes 100 percent of the switch capacity. Trying the same with a 12- or 24-fiber design either involves combining trunks with hydra assemblies or potentially stranding up to a third of the fibers at that port.

The 16-fiber building block definitely gets us to 1.6T—and, most likely, generations to follow. The other thing it does is help with power management at the switch. You can provide a lot more capacity to the network with less power per gigabit. With each new generation of switches, data capacity has doubled, but not the power needed for the switch. Energy used and cost per gigabit become more efficient over time and over generations.

And the ultra-low loss component. Can you talk to that a little bit?

As data rates increase, the industry has shifted from NRZ to PAM4 encoding, which gives us four levels of modulation on one fiber. This increases the fiber transmission capacity but, at the higher speeds, sensitivity to return loss (or back reflection) is higher for the electronics. Light reflected back to the transmitter acts as noise and can affect performance. Connections between equipment are the primary contributors to signal loss. Minimizing those losses means focusing on the fiber interfaces where any air gaps,

particulate matter or irregularities in the end-face profile can result in significant optical loss.

The end-face profile is of special concern: If it isn't manufactured, polished and cleaned accurately, the optical signal reflects back to the transmitter effecting network performance. So APC end-faces, which have long been used to control signal losses in singlemode fiber, are now beneficial for multimode as well. Transceivers using singlemode or multimode MPO16 are designed for APC, as are the trunk cables and patch cords. As demand from hyperscale and global-scale providers grows, MPO8 APC is also rapidly becoming an option. With APC, any reflected light is directed into the fiber cladding and away from transmitters—ensuring the best optical performance between devices. To minimize optical return loss and meet CommScope's stringent standards for ultra-low loss performance, Propel offers APC connector end-faces that are precision polished and assembled using CommScope's state-of-the-art fiber alignment process. Manufactured in a highly controlled and clean environment, they help ensure true ULL performance that's globally consistent and reliable.

You mentioned earlier that Propel is also designed for easier deployment and management. How?

This capability is really reflected on two levels: a design level and tactical level. On the design level, all Propel fiber trunk cabling, modules, adapters, arrays and patch cords are aligned by fiber counts. All MPO trunk cables are pinned on both ends. For MPO patch/equipment cords, both ends of the patch cords are non-pinned. This eliminates or significantly reduces the possibility of damage to transceivers from using a pinned assembly. The platform is also standardized on Method B enhanced polarity and Method B trunks. This assures transmit goes to receive and also aligns with the fiber assignments in transceivers. The same patch cords are used throughout the channel, whether at a patch panel, network device or cross-connect—simplifying your BOM design.

On the tactical level, we've designed a lot of features into Propel's panels and blades that enable faster, easier deployment, maintenance and upgrades. The panels can be installed by a single technician and then accessed from the front and rear, which is great where space is limited. Each blade has 12 rail positions that provide guides for securing different combinations of module sizes and interfaces. There's also a unique slide-out blade design that enables one technician to install and manage connectivity. Additionally, we've built in an innovative rear trunk cable manager that provides trunk cable option flexibility while holding the cables secure. One other thing to note is that the trunk cable

We've designed a lot of features into Propel's panels and blades that enable faster, easier deployment, maintenance and upgrades.

breakouts are scaled to match the shorter depth panels and reduce cable slack storage.

You mentioned application-specific design earlier. What exactly does that mean?

Propel is designed to align the infrastructure cabling with the applications. One to one. 16-fiber transceivers, 16-fiber trunk and array cables, and 16-fiber distribution modules. For 8-fiber transceiver applications, there are similarly 8-fiber trunk cables, arrays and 8-fiber distribution modules. We provide efficient building blocks to scale as you need to and manage applications without wasting fibers or losing port count. One result of this is a greatly streamlined portfolio that minimizes the number of SKUs while increasing the options and capabilities. By simplifying the portfolio and making it completely modular, everything goes together seamlessly and evolves easily.

Can you explain how the QR coding works?

All Propel modules, panels, and cable assemblies have a QR code that techs can scan with their phone and access a variety of manufacturing, testing and performance data from CommScope's WebTrak® system. The Propel QR codes are serialized and provided on modules, adapters and panels as well as with fiber trunk, patch and array assemblies, which have a QR code printed on labels located near each end. If you're testing a channel and need to reference baseline factory test data, you can use the QR code to easily access information to help validate, troubleshoot or install. That becomes a big help when timing is critical.

Will all products in the Propel portfolio be available from Day 1?

No. Due to the large scope of the Propel platform, we are releasing it in three phases.

At launch we've introduced the core products for 400G/800G applications. This includes our 144LC sliding panel in 1U, 2U and 4U sizes along with 8-, 12-, 16- and 24-fiber MPO-LC and MPO-SN modules. There are also a range of adapter packs and cable assembles included with this phase. Adapter packs feature LC-LC, MPO-MPO, as well as SN-SN adapters in the same size footprints as the four module options. MPO trunks and arrays supporting 8-, 12-, 16- and 24-fiber applications with LC Uniboot and SN cable assemblies are also included.

Next, we're adding MPO-MPO conversion modules and array assemblies, splice cassettes, ruggedized fanouts with MPO-to-LC Uniboot connectivity, and additional panel options.

Following that will be additional housing options as well as our integrated imVision® automated infrastructure management solution. So, lots to look forward to as we progress with the rollout!

COMMSCOPE®

commscope.com

Visit our website or contact your local CommScope representative for more information.

CO-116927-EN (07/22)