WDM & Next Gen Optical Networks 2011 - Monaco

Thoughts on the conference from IIR Telecoms and Technology's media partner, market analysis and marketing consulting firm Telecom Strategy Partners.

Introduction

The IIR WDM & Next-Gen Optical Networks conference in Monaco was well attended this year, and was an interesting mix of consensus and controversy. Vendors and operators alike conveyed an upbeat view on the role and prospects of optical transport as a means to help operators scale, streamline the operational model, build more dynamic and flexible networks, and support differentiated packet services as well as circuit services.

There was an unprecedented degree of sharing of networking insights and experience among operators, and vendors in turn contributed their experience and vision. Some of the presentations at this conference definitely provided "forward-looking" views on where optical technology is headed and the expected timeframe. But perhaps even more valuable was the strong practical focus that characterized the conference this year.

There was considerable discussion of what optical systems, technologies and topologies operators have already tested or deployed, lessons learned in that process, and what operators would like to see vendors supporting in the future. There was also plenty of consensus among vendors on the need for 100G and value of technologies such as coherent detection and the need for directionless and contentionless ROADMs with more degrees of wavelength switching. But that consensus was complemented with a bit more controversy when it came to discussions on topics such as the need for an OTN layer between IP/MPLS and DWDM to provide grooming in the core, whether OTN switching has as strong a role in the metro, or whether MPLS-TP is ready for prime time.

The WDM & Next Gen Optical Networking conference remains fresh because optical networking itself continues to evolve. Already we find ourselves discussing the solutions to 400G and 1 Tbps transmission; tomorrow's solutions are conceived today. But much of the focus was on current implementation of 100G and ROADM technology, planning requirements and migration strategies.

Although some people in the industry only a few years ago had predicted that transport would be little more than a commodity that operators might well choose to give away by now, nothing could be further from the truth. What *is true* is that increasing focus is placed on the dramatic evolution in new packet services and applications. But the fact remains that operators need to figure out how to efficiently provide, manage and support these new services and applications if they are to realize the associated revenue opportunities and sustain profitability. And optical networks will be a key part of that equation.

Operators need more scalable, more easily manageable, and more flexible networks than ever before. Bandwidth must be fully utilized, and planning, provisioning, and managing the network must be simpler. But more importantly, operators need more intelligent networks now. They need networks that can dynamically allocate bandwidth where it is most needed. Networks that are service aware, can provide differentiated treatment of traffic, and can provide tiered levels of restoration and protection. The need to support legacy circuit-based services and provide wavelength services has also not gone away.

With so much change, operators and vendors alike need to keep up to date on evolving opportunities, changes in technology, and the day to day proof points established and experience gained in recent trials and commercial deployments. If that is not enough reason to attend WDM & Next Generation Optical Networking, consider the substantial professional networking opportunities.

Presentation Highlights

Attendees of the conference have opportunity to access all the presentations via the IIR web site.

ADVA provided insights for operators wanting to overcome the challenges of running multi-layer management systems that create a unified end-to-end view. ADVA discussed the benefits of leveraging MPLS, GMPLS, and ASON to reduce time to services, automate dynamic bandwidth provisioning, and simplify path computation while defining differentiated protection and restoration. The advantages of integrating management, data, and control plane to enable routers to be connected to DWDM without requiring transponders, and to enable routers to use automated scripts to create loose paths across the network rather than relying on manual provisioning of DWDM equipment was a key part of the message – and supported ADVA's partnership with Juniper.

Alcatel-Lucent discussed the requirements for Terabit networking, and what it takes to achieve converged backbone transformation. Advocating backward compatibility in 100G+ electro-optics, automated photonic switching, and a sophisticated GMPLS lambda restoration with non-blocking multi-Terabit OTN switching, automated ODU networking, and sophisticated GMPLS sub wavelength restoration, Alcatel-Lucent suggested the need for coordinated multi-layer operations and a single control plane.

Belnet discussed what it takes for an operator to evolve to "zero touch" photonics based on lessons learned from the Belgian National Research and Education Network. ROADM, GMPLS/ASON, and OTN switching topics were discussed. Questioning whether OTN switching was potentially a reintroduction of unnecessary electrical backplane and additional unnecessary complexity, Belnet also questioned the value of long-term maintenance contracts with vendors in a world of dropping prices.

BT Innovate discussed its 100G trials and thoughts about future network scalability. Advocating coherent technology for the ultimate performance and reach, BT Innovate cited its very real benefits of coherent technology in terms of helping overcome the challenges of poor fiber links, enabling DCM-free network operation, simplifying network design. They also highlighted the need for 100G systems that support alien wavelengths and provide support for existing 10G wavelengths.

Ciena discussed the special challenges of 100G in submarine networks, and the benefits of building networks with long reach, that require no re-engineering or changes to 3rd-party wet plant, and are easily upgraded from 40G to 100G with a simple card swap out.

Corning discussed the continued innovation in fiber that will be needed for every part of the network as operators continue to push toward use of higher speed transmission rates. Highlighting Corning solutions for submarine, terrestrial backbone, metro and access requirements, Corning also discussed the value of low bend loss fiber for in-building use.

Deutsche Telekom discussed its 100G trials, and lessons learned regarding the challenges of deploying 100G in existing networks and of achieving 100G IP packet and DWDM OTN interoperability. Having proven the feasibility in field trials, DT discussed early experience operating 100G transport and applications

together with existing customers, and potential technology directions for the future as we move towards 1 Tbps transmission.

ECI discussed the benefits of next-gen 40G and 100G capable packet-optical transport, highlighting the value of L0/1/2/3 switching/grooming options, the need for multi-layer management and software tools. Discussing the need to support mixed 10G/40G/100G and the importance of eliminating dispersion compensated fiber in high speed networks, ECI discussed the value of router/ODU cross connect offload and ROADM/WSON by-pass and restoration.

Ekinops discussed what is needed to build scalable wireless backhaul networks that are LTE-ready. Recommending operators choose best of breed Layer 1 technology complemented by best of breed switching, Ekinops suggested that integration of layers must be approached carefully to avoid potentially sacrificing key functionality and cost efficiency, and recommended that operators think outside organizational boundaries to pick a single network solution for all their applications.

Ericsson discussed 40G, 100G, and the roadmap out to 1 Tbps. Citing ODU4 transport services, 100GE interfaces on routers, enterprise services and capacity exhaust as key drivers for 100G, Ericsson discussed the different requirements for 100G in the metro versus the core. Discussing the various means to scale network capacity and capacity per fiber, Ericsson presented some highlights from field trials and its expectations of the relative pricing of 10G, 40G, and 100G in a few years. They then went on to present a discussion of possibilities of how to attain 1 Tbps transmission rates in future.

FT Orange Labs discussed the technologies needed to move the industry to 400 Gbps and 1 Tbps transmission in future. Projecting a total capacity objective of 50 Tbps by 2020, they discussed the potential value of single carrier PM 16-QAM modulation for the metro, of multi-carrier techniques like Nyquist-WDM or multi-band OFDM for long distance applications, and of the potential use of distributed Raman amplification and/or new fibers as part of the solution.

Huawei discussed 100G technology, performance, and its experience to date with a number of customers.

Infinera discussed whether Terabit waves were destined for future use, or whether super-channels based on closely packed waves not conforming to the optical grid were more likely. Suggesting the second scenario, Infinera

highlighted reasons why it thinks its PIC technology has inherent advantages in meeting market requirements.

Juniper discussed the value of what it calls the Converged Super Core, including the relative benefits its sees for an IP/MPLS core running over DWDM as packet layer infrastructure with OTN used for transport. Suggesting that lambda switching and OTN switching have their place, *albeit a relatively small one* given the dominance of IP/MPLS traffic going forward - Juniper advocated using a hybrid switch combining MPLS switching with OTN in a Converged Super Core. This was presented as an alternative to simply feeding IP routers or OTN switches from edge routers, or to using OTN switches fed by edge routers to selectively bypass IP routers in the core.

NSN highlighted the need for maintaining simplicity while upgrading functionality on a smooth path toward packet optical convergence. They discussed the benefits of the converged packet optical network, and how the new breed of optical transport products integrate grooming and aggregation for circuit and packet. The value of MPLS-TP and benefits of ODU-level switching were discussed, along with the value of forwarding and grooming of electrical traffic onto wavelengths, and of using multiple protection and restoration mechanisms to increase network utilization and enhancing network resiliency.

The **OIF** discussed its role in enabling introduction of interoperable 100G by providing specifications and implementation agreements. Stressing that cost and energy efficiency will be key attributes of these systems along with spectral efficiency and reach, the OIF indicated its work on 100G has been designed to ensure that key concepts will be able to be re-used as transmission rates extend beyond 100G in coming years. The OIF discussing how photonic technologies such as advanced modulation formats, polarization multiplexing, and coherent transceivers have complemented electronic technologies such as DSP-based coherent detection, enhanced FEC, high speed A/D converters and common electrical interfaces to enable 100G – then discussed a number of technologies such as soft-decision FEC, new modulation formats, and enhanced amplifiers and fiber types that will enable the industry to move beyond 100G in future.

NYSE Euronext

Think your network has a lot of traffic? NYSE Euronext has to transport over 22 billion messages daily. NYSE Euronext discussed some of the special requirements of the financial services sector. First, low latency market data tends to be uncompressed, so requires more bandwidth. Second, because it

involves highly time sensitive information that can be worth huge sums of money, competitive advantage can be determined by mere differences of microseconds. If you have a “better late than never” philosophy, it does not apply here – delays can be worse than discards in the financial world.

Having tested 100G at a data center, NYSE Eurotrade has reflected on that experience and determined the need for a flat cloud network operating within the context of the data center itself, and being extensible to about 50 miles. They would like that network to support storage, bandwidth on demand, low latency and jitter, high reliability without the need for redundant switches, and the ability to scale to 10s of thousands of port of 10G, 40G and beyond. Further, they would like to see DWDM integrated into the data center router, and GMPLS fully controlled and managed from the data center routers.

Orange France Telecom discussed its transmission network evolution with ROADMs in countries in which it operates, and the value of colorless and directionless ROADM technology in helping it simplify operations, lower carbon footprint, reduce TCO, and lay the groundwork for future automation, optical restoration, and multilayer interworking. Last but not least, the role of the new network in enabling the roll out of new services was discussed.

In a separate presentation, Orange France Telecom discussed the move from 40G to 100G and beyond, and how the technical solutions required for data centers, for links between routers and WAN transport equipment, and for WAN line transmission are very different. Indicating that coherent PM-QPSK 40G and 100G WDM transmission now offer a credible alternative to 10G in LH-Metro networks without requiring a change out of existing infrastructure, Orange France Telecom suggested that Capex of less than 3.5x the cost of 10G will be required to make 40G attractive, and less than 7x the cost of 10G will be needed to make 100G attractive. They also discussed the need for 400G – 1 Tbps transmission in future, and talked about possible technology directions in terms of single and multi-carrier approaches, modulation formats, and amplification technologies needed to make that happen.

Orange Labs discussed the value and a few of the limitations of GMPLS. Suggesting that today GMPLS has mainly been applied as an “IP control plane for Optical networks,” Orange Labs suggested that the best days of multi-layer are still a bit in the future. Stressing that there are no immediate savings during a transition period but that there are very real long-term potential benefits to deploying GMPLS including dynamic restoration and automation, Orange Labs

discussed the transition toward multi-layer control and the advantages and limitations of doing so.

Telefonica Spain talked about their planning, implementation and strategy for photonic mesh. Wanting greater automation, reliability, and flexibility as well as scalability, Telefonica has seen the need for colourless/directionless ROADMs for non-blocking flexibility, photonic switching to reduce cost, a 10-degree ROADM and the ability to mix 10G and 40G. They were also particularly interested in having the automatic discovery and provisioning benefits of a GMPLS control plane, plus traffic engineering and restoration benefits, planning and engineering tools, automatic backups, and the ability to prioritize and preempt traffic on a nominal route as needed. Accounting for experience and lessons learned, Telefonica wants 20 degree ROADMs available in the future, plus integrated regeneration and lambda conversion.

Telekom Indonesia shared some of its experience in implementing OTN, MPLS and Ethernet technology in multi-vendor environments. Stressing the importance of preparing a detailed network roadmap prior to migrating the network, Telekom Indonesia indicated that they were able to use ASON/GMPLS to create a highly satisfactory implementation of a more flexible and cost effective “zero touch” transmission network – but shared that their experience with ASON was somewhat similar to IP in the amount of training required, and suggested that a deep knowledge of physical infrastructure is required to guarantee the diversity of network resources with respect to resiliency.

Transmode discussed the mobile backhaul opportunities for wholesale providers, recommending a multi-service mobile backhaul solution that supports all traffic to a cell site using a single fiber and wavelength while delivering the individual synchronization requirements of different carriers and different generations of 2G/3G/4G equipment. Discussing the needs for low latency and precise synchronization as well as support for Ethernet E-LINE, E-LAN, E-TREE and MBH per MEF 9, 14, and 22.1, Transmode suggested that EXMP provides far better support for mobile backhaul than is required by the requirements of MEF 22.1.

Verizon outlined the success of its first commercial deployment of a 100G Ethernet network for the IP backbone, describing their European backbone and testing and production deployment with key vendors and indicating that operational and error-free 100GE is already a very real possibility.

Virgin Media discussed the requirements of the network in support of the Cloud, citing capacity, data assurance, application-aware network performance, and the need for managed security services for networks.

Wind Hellas discussed the rationale for its plan for converging fixed and mobile networks; ensuring end-to-end service quality is key was a key driver. In its quest for a future-proof architecture, Wind Hellas discussed its use of hybrid microwave systems to maintain backwards compatibility and to provide an evolution path to IP & Ethernet. They decided to keep E1s in place at macro sites for synchronization, are using FE/GE for 3G and HSDPA, and are using SHDSL for Micro BTS and NodeB equipment deployed in rural areas – recommending aggregation of traffic from that split microwave prior to backhauling of that traffic over VLANs and VCGs as a simple and streamlined approach.

XO discussed its 100G trials with multiple vendors, including the internal and external drivers it faces for 100GE, and the remaining concerns it has.

Xtera highlighted the value of modular Raman amplification in supporting a mix of 10G, 40G and 100G on a 50 GHz grid to achieve up to 24 Tbps per fiber, and the benefits of better control over the per channel optical profile along the link – resulting in longer reach or higher channel counts.

ZTE discussed why it believes OTN solutions must be deployed end-to-end, and what is required to make that happen.

Telecom Strategy Partners' View

Density, scalability, manageability, availability and reduced cost per bit have *always* been important. But optical networking has increasingly become all about building *bandwidth efficient, flexible, and intelligent end-to-end networks* that make use of *every bit of capacity, provide more efficient means of protection and restoration, dynamically allocate capacity as needed, and enable differentiated packet services and SLAs.*

Optical networking is not just about supporting existing revenue streams more cost effectively. The mantra is no longer simply "lowest cost per managed bit." Optical networking has moved up the food chain to also become a critical enabler of new packet services and revenue streams.

What new challenges are optical networks addressing? Providing support for increasingly meshed traffic patterns. Enabling differentiated transport of tiered levels of video services. Supporting the capacity requirements of explosive bandwidth growth despite reduced certainty in capacity planning. Simplifying planning. Streamlining operations. Overcoming the challenges of high speed transmission over the installed base of fiber. And about building networks that can handle any service, protocol, or bandwidth requirement that the future might throw at them…while enabling the future of IT by helping enable Cloud networks and Cloud-based services, applications, and infrastructure.

Optical networking has become the science of efficiently planning and building transport and aggregation *networks*, it is no longer the fine art or complex process of fiber qualification and engineering on a link-by-link basis. Intelligent transport requires a new degree of service awareness and multi-layer management. It's not just about reducing TCO anymore…its about creating new revenue streams while supporting existing services.

What was apparent this year is that optical networking has not been commoditized by the explosion in packet traffic. By contrast, the importance of optical networking has been magnified because of the explosion in packet traffic. And the value and role of the optical network has simply increased as multi-layer integration of WDM with layers 1, 2, and 3 increases, and as optical is increasingly combined with the OTN standards-based muxing hierarchy and contribution to network interoperability, OTN switching, ROADM, GMPLS, and packet networking technologies such as VLANs to provide E-LINE, E-LAN and VPLS services and packet synchronization technologies.

100G is proven now; new higher order modulation formats and coherent detection have earned strong confidence from operators, which value the simplified network planning, distance capabilities, and wide support for a variety of fiber types that result. 100G will ramp up in volume more quickly than did 40G, largely because the pace of the escalation of bandwidth demand is even greater now than it was a few years ago. But volume dictates price, and operators still are being conservative about deploying greater volumes of 100G while waiting for price to come down. That of course requires either volume or deep pockets and a willingness to "forward discount," knowing full well other vendors will not follow that pricing strategy and would price aggressively in future - making it difficult to recover target levels of margin). So 100G has a pretty solid showing in terms of performance, but remains a small part of the market so far in 1H 2011. But that will change.

40G will still have a couple of years of good opportunity overall, and potential extended good opportunity beyond that for a time in selected parts of the metro. CapEx may not be low enough yet to induce wide scale deployment of 100G, but operators are leveraging 100G where needed, gaining experience, and actively looking forward to the future. Single-carrier solutions will be the way forward at 100G now due to power and spectrum savings, but at 400G and above the jury is still out, as one way to get there will be using a multi-carrier strategy. 100G will of course ramp more quickly in the core, but its need in the metro will not be far behind – though the technologies needed to address the metro will be different.

Coherent detection is a great example of a technology that increases cost, but is very desirable to operators nonetheless because it adds so much value that the cost/benefit ratio is quite compelling. The once “insurmountable” problems of chromatic dispersion (and also PMD) at 100G have effectively been overcome. Having already gained wide acceptance, this is a technology that has not only had a huge positive impact, but will be with us well into the future.

ROADMs have become table stakes on almost any RFP these days, and are a key technology needed to make networks more flexible and efficient, allocating capacity where needed. They’re also an important building block needed to support the creation of a flexible, intelligent, automated network leveraging a GMPLS control plane. Operator interest in colorless and directionless ROADM technology is particular high, and there is strong interest in support for as many as 16-20 directions in future. Vendors have also more recently been touting the advantage of contentionless ROADMs more recently, but that seems to have attracted a bit less attention from operators, at least in the short term.

Optical networks provide a key underlying transport layer, but that transport layer increasingly will need to be intelligent, automated, and aware of higher network layers – providing differentiated treatment of traffic for various services and SLAs. GMPLS was developed a good decade ago, but for many operators, the topic is becoming much more important now. The full benefits of GMPLS will result in intelligent, automated networks that make more efficient use of bandwidth, effectively allocate capacity where needed, enables restoration, and ultimately helps streamline operations costs. Increasing coordination between multiple layers in networks, with an effective end-to-end multi-layer management view, will be important in future.

The need for cost effective use of all network capacity has been getting operators to increasingly move away from rings toward more meshed restoration schemes – though not completely, with protection methods often remaining in hybrid mode. Having multiple levels of restoration at different network layers also provides flexibility to enable a lot of new differentiated services.

Adoption of some of the early generations of packet optical transport platforms has been a bit slower than some vendors hoped, but as the technologies involved mature and standards solidify, there is a stronger rationale for packet-optical transport than ever. Convergence of networks is something only a few operators have fully moved to do in the short-term, but no wonder – that is a major task and will not be the first one that most operators feel the need to take on during their transformation. But many operators are already upgrading and implementing their networks in such a way that leaves the door open to the possibility of converging all traffic onto one network in the future.

The controversy between IP/DWDM advocates and those supporting an intermediate OTN layer is an interesting one. We think IP/DWDM will have its successes, particularly in data centers and research networks and some green field opportunities. But OTN switching with ODU-level granularity will prove very valuable for many established service operators, at least in the foreseeable future. 100G is only going to be cost effective if wavelengths are efficiently groomed to make full use of all that bandwidth, there are real advantages to having protection of fully packed pipes at the optical layer, and the strong legacy service support and transparency characteristics of OTN framing and familiar transport OAM model and hierarchy make sense.

The future may be packet-centric from a services standpoint, but that does not mean that all services will need to be routed, or that services do not need to be supported by a strong OAM model in an underlying transport layer. Technology paradigm shifts occasionally outdate maxims, but we believe that the “switch where you can, route where you must” maxim still has a lot of validity left in it from a cost standpoint and may well have it well into the future. There will still be a need for L2 switching, not to mention continued support for a not insignificant amount of legacy circuit traffic.

We think Mervyn Kelly, EMEA Marketing Director for Ciena said it best during our interview with Ciena at the conference: “Use IP routers *when* you need them. But that’s an expensive way to build a network for all your traffic needs. Having an inexpensive L2 blade in an expensive box doesn’t solve the cost problem,

particularly not when 75% of the traffic entering the network does not need to be routed at L3 in the first place. Use of L3 has to be context-sensitive, used only at those nodes where it is really needed."

The ROADM and OTN switching combined are going to provide significant flexibility in the core, and are key to setting the stage for meaningful implementation of GMPLS. In the words of Sam Bucci, VP of the Terrestrial Optics Product Unit at Alcatel-Lucent, "…customers can really benefit from the integration of tunable ROADM and multi-terabit OTN switching, offering an effective tool-set for improving overall bandwidth management and efficiency, particularly in the core of the network. GMPLS-based control plane, available both at the photonic and electronic layers, adds the required network control and intelligence to Optical networks, with the goal of providing the right level of robustness, resiliency and network automation."

The benefits of OTN switching in the core are clear, but it is not necessarily the right technology everywhere. Along those lines, David O'Neil, Senior Product Manager Transport, ECI Telecom, made relevant comments we agree with when he said "In smaller networks utilizing ROADM protection/restoration might be the most cost effective solution and would save I/O ports, but in larger regional or long haul networks where regeneration of back-up paths is necessary an OTN ODU cross connect fills this function as well as letting you change wavelengths to better utilize the network, and offers the advantage of providing efficient grooming of a variety of traffic ranging from 10G, 40G, or 4G or 8G Fibre Channel to let you very efficiently and completely utilize your 100 Gbps wavelengths and allowing you to maintain end to end visibility at the service layer. You want to protect and restore traffic at the lowest possible network layer to reduce cost; we have to look at that on a case by case basis to determine the best solution for each operator's network and for each part of that network. Converged Packet Optical transport Systems with Layer 0/1/2/3 switching . grooming and protection/restoration options can provide the most versatile and cost-effective solution for any operator requirements."

The optical world has been rapidly changing, becoming intertwined with packet technologies and higher layer networks. That will continue. As Sten Nordell, CTO of Transmode observed, "…It is really interesting to see how quickly the optical networking industry is evolving to support new requirements such as Sync E support within packet-optical solutions. Furthermore now that we have these solutions rolling out, we are seeing that our customers can provide better quality synchronisation with SyncE than the previous PDH/SDH/SONET networks, along

with the better economics and scalability of packet-optical."

The world is becoming more complex, and that gives certain advantages to large vendors. But note also the partnership between Juniper and ADVA – integration of the routing and DWDM layers does not necessarily have to be a single vendor effort, and there is room for innovation from players with more focus as well.

The former trend to outsource all component manufacturing is now officially over; that was a “sign of the times” when the market was slow and infrastructure was commoditizing. Vendors are clearly turning once again to strong innovations in silicon as a means to achieve differentiation, both in optical networking infrastructure and in switching, routing, and packet optical transport portfolios. Advances in silicon are critical to enabling digital signal processing for coherent detection in optical, and can scale differentiated treatment of traffic on a per service, port, or subscriber level in accordance with policy.

Ericsson’s and France Telecom’s presentations were among those clearly illustrating the amount of thought already going into surpassing 100G technically in the potentially not too distant future – though we think the ramp up following commercialization of that will not occur before the end of the decade.

What is Driving All The Change

The need to support new service opportunities with differentiated performance and radical new network is the key driver of change in optical networking. Vendors are responding with high performance transmission at increasing line rates over any fiber type, the ability to mix channels of various transmission rates over the same band, reduced need for dispersion compensating fiber, planning complexity, meshed restoration, reduced power requirements, and increased manageability.

Optical transmission continues to offer huge advantages in both scalability and cost per bit. Bandwidth requirements are unprecedented, continuing on their spiral of rapid growth, and are increasingly hard to predict. New service revenue opportunities can be profitably realized only if a whole new level of cost optimization can be achieved, due to the lower cost per bit that many of the new services generate.

These changes requirements are driven by the increasing sophistication of the communications devices subscribers use, the increasing mobility, graphics

capabilities, and processing power of those devices, and the seemingly insatiable demand for video and social networking in the mobile as well as fixed context.

The advent of tablets and small, lightweight, high performance personal computing devices with USB modems is driving mobile traffic requirements to unprecedented levels. This has been encouraged by the mobile broadband capabilities offered by HSPA+ and LTE. Subscribers are carrying multiple communications devices, each capable of higher download rates than in the past. And they are engaging in an unprecedented degree of peer-to-peer communication.

Machine to machine communication will cause an explosion in the number of network end ponts. There are also a host of new bandwidth-hungry applications and services requiring differentiated performance, and video content has begun to dominate network traffic. Cloud computing promises a network-led revolution in the way we leverage IT software, platforms, and infrastructure, and the IT requirements in turn are changing network requirements as well.

Network scalability requirements have simply exploded. Peak-to-average traffic ratios have blossomed, traffic patterns are becoming more complex than ever, and capacity planning has become extremely difficult if not impossible. More efficient aggregation, forwarding, and transport of an increasingly packet-centric mix of traffic is required. Every bit of available network capacity must be efficiently used, causing traditional ring and 1+1 protection schemes to increasingly give ground to a greater degree of meshed or networking and restoration – with hybrid schemes becoming increasingly common.

QoS requirements are more granular than ever, and there is a need for differentiated treatment of traffic at the per flow and per subscriber as well as port level. Packet-based services create opportunities for new tiered services offering multiple levels of service differentiation and oversubscription. And the need to support and report on SLAs is critical, both as a means to guarantee quality of experience for subscribers based upon operators' internal network objectives and KPIs, but also to to provide verification of network performance and availability as an explicit service in itself.

Regeneration must be reduced, and footprint and overall power requirements minimized. Dynamic allocation of network resources is highly desirable, and provisioning must be streamlined in terms of both time and cost. Every bit of

potential bandwidth carrying capacity must be drawn from the installed base of fiber, and innovation in fiber types also continues.

Innovation in services is driving the need for innovation in optical networking, and vendors are responding. Though we expect greater integration of optical with other technologies in single devices to continue as a trend, it is not because optical has become a commodity or has become less important. It is because optical networking has become an invaluable part of solutions for which additional synergies in packaging and reduced cabling requirements are needed to reduce cost, and for which there will be growing synergies between optical layer and L1/L2/L3 technologies.