WHITE PAPER

Evolution of Network Design

The impact of networking trends on network design, construction, and operation and what that means for networks moving forward.

Network bandwidth demand has grown relentlessly over the last 20 years. Internet and mobile networks are experiencing unprecedented traffic increases, with no sign of slowing down. As a result, networking designs have had to shift to accommodate these requirements—adding complexity and operational pressures.

These networking trends include:

- Moving fundamental network building blocks from SONET/SDH to Optical Transport Network (OTN) and Dense Wavelength Division Multiplexing (DWDM) technologies
- A shift in networking to predominantly IP-based traffic traveling over Ethernet, with metro and long-haul Ethernet traveling over primarily OTN or DWDM
- The need for IP/Optical convergence to increase network efficiency
- Employing sophisticated software tools for network provisioning and management
- Leveraging virtualization (virtual machines or containers) of network functions to increase the flexibility and agility of what traditionally were hardware-based appliances

Further, software control has become both necessary and desirable to operate networks. As network complexity has increased, sophisticated software tools are required, and operators demand software that offers further flexibility and programmability to best utilize their networking assets.

The combination of compute, storage, and connect has enabled the rise of the cloud, comprised of public, private, and hybrid cloud architectures. The cloud is physically housed by both large and small data centers that are both centralized and highly distributed. Software is now being designed as 'cloud enabled' or 'cloud native,' which further increases traffic across networks.

All these networking trends not only increase the demand placed on networks, but also make network design, construction, and operation much more complex. Moreover, current and projected networking trends show an increased demand in the future:

- The transition from 4G to 5G has begun, with 5G supporting the increased bandwidth, density, and low latency demands of end-users. Work is already underway to define 6G, which although not formally specified yet, will likely include an exponential increase in the performance demanded from the network.
- Quantum computing, including quantum key encryption and quantum key distribution, looks to not only increase the computing power of the cloud but also increase the security of the data running through or stored in it.
- Virtual reality, extended reality, and augmented reality promise applications that require huge amounts of network capacity and very low latencies—or both simultaneously.
- IoT holds the promise of a massively connected, fully instrumented, and extremely smart/intelligent world.
- Artificial Intelligence (AI) and Machine Learning (ML) will likely be needed to operate all the networks and systems without humans directly involved.

Changing traffic patterns

ICPs drive new services and set the technology agenda

Technology evolution and direction is increasingly driven by Internet Content Providers (ICPs) such as Amazon, Facebook, Alphabet, Apple, Microsoft, Tencent, Alibaba, and more. This is true for both Wide Area Network (WAN) capital expenditures, as well as inside data center technologies such as low-cost, low power network edge solutions that are driving data volumes and pushing capabilities into the WAN.

First, ICPs lead traditional service providers in spending on network equipment. As such, network designs are starting to show much more attention to the specific requirements driven by ICP models, with traditional service providers tacitly acknowledging this transition.

Second, ICPs are driving much higher volumes into the networking segment inside data centers. This disparity will increase as the inside data center growth rate is higher than WAN, driven by the service mix. AI and ML algorithms are likely to exacerbate these trends even more, as they demand huge internal interconnect bandwidth for learning and data processing (like Google's special Torus network for tensor processor nodes). The approaches for electronics, optics, and software designs that have been perfected for inside data centers are being extended into the access and metro space. These trends produce a traffic bandwidth distribution where intra–data center traffic is dominant, which is then followed by data center–to–user traffic, and finally with data center–to–data center traffic. As data center single-building size hits a limit and multi-building campus deployments proliferate, a larger fraction of traffic that is indicated as 'within data center' will interconnect campus buildings and will leverage technologies designed for data center–to–data center interconnections.

Access and aggregation architectural evolution

As discussed previously, the industry is at the precipice of the next step in growth. Bandwidth has been the primary driver in previous evolutions of network architecture and continues to be critical, but the next step will be more than just additional bandwidth. Not only will billions of endpoints enabled by 5G create a massive influx of data onto the network, but where content is stored and how it is accessed will fundamentally change. Content will be distributed further toward the edge of the network for better Quality of Experience (QoE), for example, lower latency and better reliability. Traffic patterns will evolve from an 'access-to-core' paradigm (such as, point-to-point) to an 'any-to-any' distributed model, again shifting emphasis toward the network edge. Traffic will become less asymmetric as users create video content as well as consume it. More recent evidence for the acceleration of this trend is provided by many large content owners building their own edge-oriented delivery networks.

This requires a rethink of the traditional 'hierarchical' network architecture, as the above factors will drive the requirement for capacity closer to the edge of the network. Historically, higher bandwidth demands would require separate DWDM transport systems, which in turn would demand more space and power; this is often unrealistic in many locations due to physical constraints of the environment.

The industry is evolving toward miniaturized modems, enabling new options to address these pressure points. Several new form factors are emerging which may yield interesting options for operators when configuring their access, aggregation, and transport domains. Platforms are emerging that combine high-capacity routing and switching, highly dense access interfaces, and coherent optics. Platforms integrating high-density IP/Ethernet routing and switching with coherent modems enable packet aggregation with high-capacity transmission in a small power and space footprint that can be deployed further into the access network.

Transport networks

Dynamic optical layer

Network demand is constantly growing, and a static network mode of operation is no longer sustainable. At its core, Layer 0 (L0) dynamic adaptation exploits the ability of optical hardware to precisely match the capacity of an L0 channel (or wavelength) to the system margin—in contrast to the attempt at higher layers to optimize the capacity of a mesh network. The 'wavelength channel capacity' is determined by Signal-to-Noise Ratio (SNR), and if SNR is known, network performance capacity and unregenerated reach can be optimized to achieve CAPEX and OPEX savings.

More generally, three capabilities would be required to operationalize a fully dynamic optical layer:

1. An electrical Layer 2 (L2) or Layer 3 (L3) capable of delivering variable client bandwidth to the optical channels, programmable optical hardware, and intelligent software.
2. A flexible grid and a reconfigurable photonic layer with a colorless, directionless, and contentionless capability to allow for rerouting channels of variable spectral occupancy across any path.
3. Variable bitrate and software-configurable coherent optics which can match optimal channel capacity to available SNR for a specific network path; protocols such as Flex Ethernet (FlexE) can efficiently map a flexible number of client signals to the variable line capacity as needed.

There are several external constraints imposed on the network dynamics and associated optimization problem, which can be grouped into two broad categories:

First, it is attractive from an optical line-system perspective to focus on optical L0 adaptation. The simplest approach is to keep modem Baud rate—and thereby spectral occupancy—fixed and adapt the modulation format to fit the required network optical reach. This is easy to implement in a modem and simplifies line-system planning as spectral occupancy stays constant. However, this leads to possibly substantial changes in total bandwidth throughput and necessitates large changes in bandwidth through connected electrical-layer switch fabrics and interfaces. Switch fabrics and interfaces leading to every optical modem must be designed for maximum possible bandwidth—but may actually operate at lower utilization, leading to effectively stranded electrical-layer hardware capacity.

Second, as most traffic originates and terminates as an IP service, it is natural to consider electronic routing as the dominant layer. The interface bandwidth between a router and the optical transport layer is fixed—either by the internal electrical-switch capacity or external client interface (for example, 100GbE plug). The problem of electrical hardware utilization and total network bandwidth throughput is simplified and stays constant. However, this moves optimization complexity into the optical layer. Modems must now change both Baud rate and modulation format to keep overall throughput constant. This presents an issue, as modem cost and power are driven by the Baud rate which sets the Radio Frequency (RF) bandwidth of modem electro-optical components and Analog-to-Digital and Digital-to-Analog Converter (ADC/DAC) sampling rates. Underutilizing modem Baud rate is difficult and costly—as very granular rate changes must be supported and wasted RF bandwidth is expensive. Further, the line system must now contend with variable spectral occupancy, again with very fine-grain resolution and ability to preserve spectral continuity for each channel across the full optical reach.

Advanced software applications will abstract complexity associated with flexible technologies, enabling operators to fully realize the benefits associated with the modernized network. The marketing is exciting and may indeed entice customers to demand and pay for perceived functionality. However, the actual utility of a dynamic optical layer is less certain:

- First, the network bandwidth growth rate is still around 30 percent, and so whatever suboptimal capacity is left unfilled in the network will likely not persist for long; dynamic optimization is valuable for networks with high churn and slow growth—not the other way around.
- Second, the impact of high diurnal traffic variability complicates the evaluation of network load and optical margin tradeoffs.
- Third, the optical layer is already quite efficient.
- Fourth, since the optical layer is fundamentally 'analog' and its reconfiguration is slow, it cannot be used as a direct substitute for adaptability at higher layers.
- Fifth, CAPEX and OPEX optimization typically produce configurations that fill optical transceivers to their fullest capacity (and leave little opportunity for reconfiguration).
- Finally, L0 adaptability requires that bookended transceivers be fully compatible and of the same vintage; this may slow down the introduction of new transceiver technology, especially considering the added complexity of L0 control software. The industry is working to provide optical interoperability for shorter reaches, but interoperability for high performance long-haul links is not currently being considered.

Therefore, it is likely that optical L0 adaptability may find some applications in access and metro networks, where customer and traffic churn occur most frequently and optical link performance is less demanding, but less likely in long-haul and submarine networks.

The problems of dynamic networks described above are further exacerbated by the large number of amplifiers, Reconfigurable Optical Add/Drop Multiplexer (ROADM) configurations, and transponder data rates, which lead to a huge network modeling and configuration verification load for equipment suppliers and operators. This opens an opportunity to consider an alternate approach: combining low-cost, low-power transceivers optimized for a single span reach, along with low-power deterministic electronic switching.

Advances in Complementary Metal Oxide Semiconductor (CMOS) technology are starting to produce purely electronic switching that is competitive with ROADMs in capacity and is substantially more cost-effective. A Reconfigurable Electronic ADM (READM) layer may be viewed as a direct substitute for a ROADM, but with much greater capability. For example, consider a hybrid system that combines the best of both electronic switching and optical bypass. A purely electronic transport layer could leverage an electronic 'digital' control plane, and potentially provide much faster and cheaper system development and much faster operator deployment. Planned industry interoperability of coherent pluggable modules removes operator concerns about long-term system lock-in and supply-chain risks.

IP networks

IP is the universal language of the internet—the glue holding everything together. Historically, network design was centered around IP: applications were limited to what was supported by existing IP capabilities, and evolution meant more protocols added to the existing IP software stack.

During the first three decades of the internet, there were no significant changes to this approach. It was a highly competitive game in which IP equipment vendors dictated a provider's ability to support new applications within closed and proprietary protocols. In doing so, these vendors controlled the market. The consequences of this model were quite clear: a slow innovation pace, high infrastructure refresh rate, vendor lock-in, limited supply chain, limited choice, and fast-growing operational cost and complexity for the network provider.

Within the last decade, the balance of forces has started to change. ICPs have emerged in force, disrupting the traditional notion of a network with different views on the existing network ecosystem. ICPs have no predefined bias about how to design, deploy, or manage their networks; they are focused on creating the most efficient content delivery mechanism that leverages the most advanced, best-in-breed connectivity technologies in conjunction with storage and computing capacity. In addition, end-customers' perception of value has shifted from connectivity to overall QoE. This has increased pressure on incumbent service providers to deliver a faster time to market and QoE at a much lower cost.

These changes require IP-based flexibility to move closer to the network edge to reduce transport cost and improve performance, which will result in a significant increase in the number of IP nodes. The traditional IP-centric way of building networks is simply unsustainable. Associated revenues no longer support the cost of constant capacity and platform upgrades. The ever-increasing inefficiency of continually adding more IP protocols makes network operations overly complex and unmanageable. A dense, siloed IP infrastructure has simply become too big an obstacle for operators to cost-effectively scale to new and emerging demands.

Service providers and enterprises already face an enormous challenge to support current services and applications while cost-effectively addressing ever-increasing user demands. As the speed of innovation accelerates, new technologies such as 5G, IoT, Edge Cloud, and AI create different network requirements that result in increased operational complexity with a direct negative impact on OPEX.

In legacy IP architectures, each IP platform needs a full stack of IP protocols to handle different applications. The IP platform also needs to interact with many different nodes to identify an optimized route to deliver the content. This box-centric approach is extremely inefficient, as the platform wastes a lot of capacity processing outdated protocols and signaling to a large number of nodes.

This monolithic approach makes performance and ability to scale difficult—the IP platform makes all routing decisions with a limited view of the network and the application requirements. Legacy IP routing decisions are usually far from optimal; any new application may require massive, networkwide software and hardware upgrades to support it.

The new IP network must be open, programmable, disaggregated, and virtualized in a way that allows resources to be reconfigured rapidly and without physical intervention to enable both existing and emerging services. It must support open, standards-based APIs such as NETCONF/YANG and provide rich telemetry for introducing software-defined control to self-diagnose, self-optimize, and self-heal. It must also allow for intelligent, data-driven, intent-based automation and an increased level of dynamic decision-making and subsequent actions.

Network evolution is not about adding more protocols or upgrading IP boxes. It is about learning how to most efficiently connect users to content, and the speed at which the network can adapt to new application requirements. It is about optimizing flexibility, cost-efficiency, and performance—creating a network with adaptive capabilities.

Visit the Ciena Community
Get answers to your questions

Ciena may make changes at any time to the products or specifications contained herein without notice. Ciena and the Ciena Logo are trademarks or registered trademarks of Ciena Corporation in the U.S. and other countries. A complete list of Ciena's trademarks is available at www.ciena.com. Third-party trademarks are the property of their respective owners and do not imply a partnership between Ciena and any other company. Copyright © 2022 Ciena® Corporation. All rights reserved. WP248 6.2022