Message from the Series Editor

Delivering on the 100GbE Promise

Milorad Cvijetic, NEC Corporation of America
Peter Magill, AT&T

Milorad Cvijetic

Peter Magill

In today's environment of converging networks, a transition to a unified, packet-based, core/edge-network architecture is occurring. The core network will have a two-layered architecture with packets sent over DWDM (Dense Wavelength Division Multiplexed) transport. This is the most cost-effective way of providing shared network capacity, of unifying protection schemes, and of delivering guaranteed end-to-end performance. At the same time, the edge network becomes a place for the manipulation of both data streams and services, through traffic grooming and switching. Finally, the access network facilitates the distribution of services to individual end-users, such as corporate data centers, government departments, medical facilities, entertainment organizations, scientific institutions, as well as private residential users.

Presently, most Information Technology (IT) services, and certainly all *new* IT services, are packet-based. In particular, in the residential sector, IP-television (IPTV) is a rapidly-growing, high-bandwidth service. The various types of IPTV can be categorized as

a) Streaming of high-quality content selected from a few hundred "channels,"
b) Video-on-demand (VoD) with material from a few thousand professionally-produced, stored movies and programs
c) User-generated content like that from Web 2.0 and social networking sites.

Delivery type (a) makes heavy use of multi-cast (i.e., a one-to-many protocol from each channel source to all viewers of that channel, at that time) while, in contrast, (b) and (c) require unicast (with one-to-one sessions between server and each user). Each of these types of IPTV today only requires large bandwidth (40 or 100 Gb/s) at a few places in the network. But, given that the number of television shows which actually requires real-time streaming (e.g., sports and other live events) is small, as VoD services grow, (b) will rapidly demand more network resources than (a). Then, if the video quality available from the newer applications (c) becomes comparable to that of VoD, then the network requirements of (b) and (c) become indistinguishable (ignoring business models). The transition from type (a) IPTV to types (b) and (c) will cause an increase in both the number of content sources and the number of different "channels" to be delivered, rendering multi-cast less effective.

Recall that the industry is currently increasing residential broadband IP rates (driven by the emerging IPTV services of type (a) and (b)) from 1–6 Mb/s to 25–30 Mb/s (via VDSL over copper) or to ~80 Mb/s (via fiber-to-the-home). That "bump" in access rates, combined with the greater diversity in the content being delivered at any one time, will eventually result in an order-of-magnitude increase in aggregate demand for the Internet Service Providers (ISPs). ISP backbone links today are being installed with bandwidths of 10 Gb/s. Therefore the need for 100 Gb/s rates is not far off. Furthermore, the need for 100 Gb/s Ethernet is clearly established because

- This increased traffic is all IP-based
- Nearly all these IP packets are encapsulated within Ethernet frames somewhere along their journey from origin to destination
- The need for the classic TDM channelization features provided by SONET/SDH have run their course
- Ethernet encapsulation is simpler and less expensive than SONET/SDH encapsulation

For all IT users, residential and enterprise, there is an amazingly rich mix of services available from the search and web-application providers (e.g., Google, Yahoo, Windows Live …). These web-based applications include many productivity and collaboration services such as mapping/directions, auctions, retail sales, stock tracking, games, shared calendars, photo sharing, editors and many others. While none of these requires much bandwidth at the client (by design), the aggregate communication load at the server farms exceeds 10 Gb/s today. And with a growing user base, the links both within the server clusters and from them to the network would be more efficient with an Ethernet standard >10 Gb/s.

Of course in addition to these web-apps, traditional Internet access is a critical service for all IT users today. And while the number of Internet users grows daily, so too does the average computing power and I/O bandwidth for each of them. The net result is a very large traffic load on the core IP routers and the network interconnecting them. These core routers are more efficient and cost-effective when built with the highest port speeds available. The efficiency arises because today's routers duplicate most protocol and route processing on each interface card, rather than centrally. Therefore, fewer interfaces usually translate to more efficiency. And while the per-port costs can be quite high, the cost per Gb/s goes down as the common equipment is shared and the port density increases. So today the core routers are being rapidly deployed with 40 Gb/s ports. In a couple of more years, it is clear: 100GbE will be needed for this purpose.

In addition to the core routers, edge (or access) routers are being pushed to higher and higher capacities. In the near future these routers will be equipped with 40 Gb/s interfaces. And, for the same reasons as above, these ports would be more cost-effective with Ethernet interfaces >10 Gb/s. And for many years, Ethernet interfaces have been priced much lower than their SONET/SDH counterparts. This trend is expected to continue because demand for very high-speed SONET/SDH packet interfaces is generated only from long-distance carriers. In contrast, the demand for Ethernet interfaces is shared among the much larger metro and enterprise community and therefore will enjoy more volume-driven cost reductions.

And finally, enterprise (i.e., business) users have their own service needs (in addition to those above). These customers have a growing use for virtual private networks (VPN), so that their business communications and transactions can be accomplished in complete security. Depending upon the size of the company and the number of its customers, these communications may be carried over either Ethernet-VPN or Ethernet private line. Here again, the traffic load at the customers' locations (into the access portion of the network) is in the range of 10 Mb/s to 1 Gb/s. However, aggregated up in the carrier's network, the demand on the Layer-2 switches is growing rapidly. Soon enough, the connections to and between these switches will also require more than 10GbE.

This brief overview clearly shows that the number of services requiring packetized 10 Gb/s bandwidth pipes is increasing rapidly and that the carriers need their aggregation and transport to be more efficient. There is wide consensus today on a need for 100 Gb/s Ethernet, or rather, comprehensive network solutions that run on a 100 Gb/s Ethernet (100GbE) infrastructure. That includes development of server interconnects, definition of the MAC layer parameters and LAN physical interface specifications for the 100GbE PHY, and definition of transport parameters that would allow communications over 100GbE across the Wide Area Networks (WAN) and nationwide networks.

The 100 Gb/s bit rate is a logical extension of the Ethernet hierarchy (10M/100M/1G/10G), which has been very successful over the past decades. At the same time, there would be value in developing a server-interconnect solution at 40 Gb/s, which is driving the 40 Gb/s-Ethernet proposals introduced recently. Some technologists believe that 100GbE is already late, and they have been working with vendors who have proprietary 100 Gb/s Ethernet solutions. This indicates that standards for 100GbE should be finished as soon as possible.

Two leading standards bodies, the IEEE 802.3 Higher Speed Study Group (HSSG), and the International Telecommunications Union (ITU-T), have recently taken important steps towards defining and standardizing different aspects of 100GbE. It should be noted that this effort will also be defining a 40GbE specification, targeted at server connectivity in datacenters. There are several other alliances and associations, such as the Alliance for Telecommunication Industry Solutions (ATIS), the Optical Internetworking Forum (OIF), the Ethernet Alliance, the Road to 100G Alliance, and the Optoelectronics Industry Development Association (OIDA) that are also involved in the process of delivering a 100GbE standard.

The purpose of this special issue of IEEE Communication Magazine is to explore different aspects of 100GbE with its readers. It should serve to further explain the market drivers for 100GbE introduction, current technical solutions for client and line interfaces, progress toward standardization of 100GbE, and key component technologies needed for 100GbE applications.

There are six articles included in this special issue, covering the topics mentioned above. The articles were written per invitation by well-known experts in the area, who actively contribute to the development and progress of high-speed networking and transport technologies. They come from different segments of the industry (system vendors, component vendors and solution providers).

It was our pleasure, as guest editors, to work with the authors of the articles that follow. We would like to express our appreciation to each and every one of them for their efforts to articulate the relevant topics and issues, and bring those issues to the readers of IEEE Communication Magazine.

BIOGRAPHIES

MILORAD CVIJETIC (Milorad.Cvijetic@necam.com) currently serves as Vice President and the Chief Technology Strategist for NEC Corporation of America, Optical Network Systems Division, in Herndon Virginia. He received his Ph. D. in Electrical Engineering in 1984. Having joined NEC in 1997, he oversees new technologies and product introduction for Next Generation Networks. Previously he has been with Bell Northern Research — BNR (later NORTEL Technologies) in Ottawa, Canada. He published more than 70 technical papers and three books titled "Digital Optical Communications," "Coherent and Nonlinear Lightwave Communications," and "Optical Transmission Systems Engineering." He has also taken part in numerous technical conferences and symposiums, in some as a conference/session chairman, technical committee member, short course instructor, or invited speaker.

PETER MAGILL received his B.S. in Physics from the University of Dayton, Ohio in 1979 and his Ph.D. in Physics from the Massachusetts Institute of Technology in 1987. He joined AT&T Bell Labs a month later; working at the Crawford Hill Lab on the characterization of advanced lasers, optical access networks and data-over-cable access protocols. He then went with Lucent Technologies as it was spun out of AT&T in 1996, to head Bell Labs' Access Research department. He managed the R&D of passive optical network (PON) systems and cable modem headend equipment. In 2001 he returned to AT&T and is now Executive Director, Optical Systems Research in Middletown, NJ concerned with advancing fiber communication technologies for the entire network (inter-city, metro and access) including 100 Gb/s transmission systems and dynamic wavelength services.