100 Gigabit Ethernet is Here!

Introduction

Ethernet technology has come a long way since its humble beginning in 1973 at Xerox PARC. With each subsequent iteration, there has been a lag between time of standardization and large scale adoption. The latest iteration, dubbed 802.3ba by the IEEE Higher Speed Study Group (HSSG), was ratified in June, 2010 and follows this same pattern but with a slight twist. For the first time in Ethernet history a single standard defines two separate speeds; 100 Gigabit Ethernet (100GbE) as well as 40 Gigabit Ethernet (40GbE).

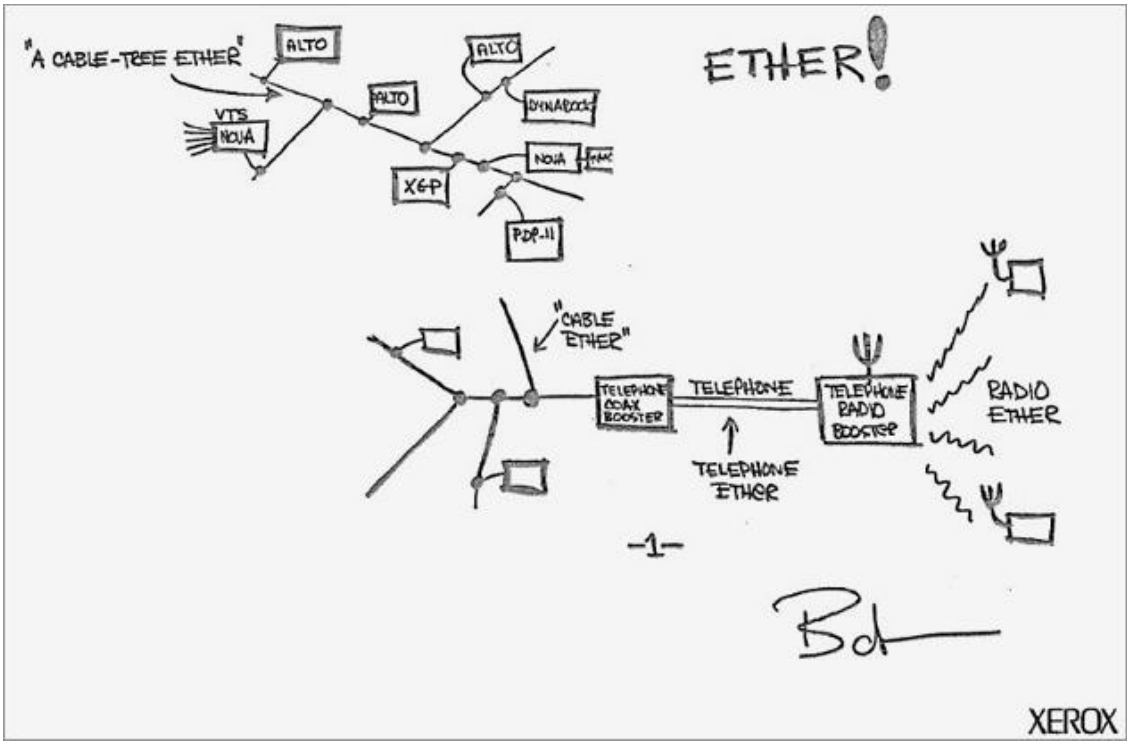

Figure 1: Original Ethernet Sketch

The technical challenges facing 100GbE have been significant; ranging from developing a whole new generation of optics that can handle 4 lanes of 25Gbps, to simply dealing with normal router and switch functions such as packet inspection, queuing, lookups, filtering and table updates, all in one-tenth the amount of time as with 10GbE. And of course this all has to be done with complete backwards compatibility and meeting all expectations with respect to bit error rate, latency, jitter and the like. As expected 40GbE gained some level of market acceptance first, but some 5 years after ratification the time for 100 Gigabit Ethernet is now!

This whitepaper will discuss the evolution of 100GbE technology in the service provider and data center markets and provide insights in to how network application acceleration hardware can be leveraged to maximize performance and efficiency in emerging 100GbE network appliances.

100GbE in Service Providers Networks

100GbE is rapidly approaching large scale adoption in the wide area network (WAN), which is largely the purview of service providers. This technology is often referred to as "Carrier Ethernet" because there are extensions needed to enable telecommunications service providers to utilize Ethernet technology in their networks. For example integration of 100GbE signals in to the Optical Transport Network (OTN), which is a standardized framework for deploying optical networks that can transport various types of data signals (including Ethernet) over long distances using wavelength-division multiplexing (WDM) technology.

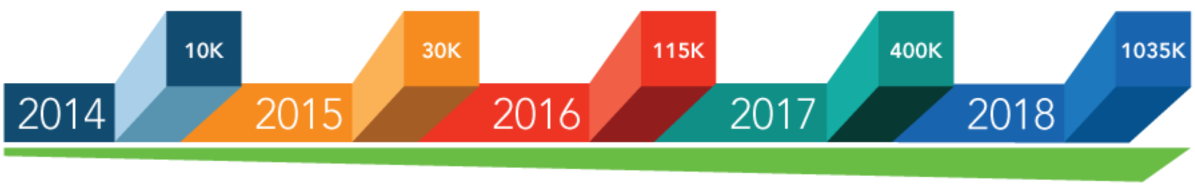

Figure 2: 100G Switch Ports Forecast

Service providers were the first to adopt 100GbE, and the earliest standards compliant deployments were in 2011 on long-haul, Internet backbone links between major cities in Europe and the US. These links, which are often 500 miles or more, connect core routers from vendors such as Cisco or Juniper via optical fiber and constitute the inner sanctum of the Internet. Because these links are at the Internet's epicenter, as traffic volumes grow due to the latest viral YouTube video or increased use of cloud services by corporations, these links are the first to saturate. Service providers must stay ahead of demand by deploying new technology and at the same time increase efficiency. Prior to 100GbE, most long-haul links were 10G with some 40G here and there mostly from larger backbone providers. Consolidating network traffic onto a single 100G link rather than multiple 10G links is cost effective because it allows backbone providers to easily ramp up capacity, better utilize fiber capacity and overall is less error-prone.

Starting from the core of the Internet in 2011, 100GbE has steadily moved outward in service provider networks. The next stop was large aggregation networks which feed in to the core and then finally towards metro area networks (MANs) which use Ethernet technology to provide high speed Internet access and facilitate communication between business, academic and government entities in a metropolitan area. 100GbE is now becoming well established in service provider networks and deployments are continuing to grow steadily.

100GbE in the Data Center

The first production, enterprise class 100GbE switch ports emerged in 2013 for use in large-scale data centers that support application environments such as for high performance computing (HPC) and storage.

100GbE is emerging at a time of fundamental change in data center architectures that is driven mainly by the advent of virtualization and associated large scale server-to-server communication. The core, aggregation and access three tier architecture was common when data centers were largely designed to support remote client-server applications such as email, ERP or collaboration portals like Microsoft Sharepoint. In this architecture, traffic was mostly traveling “north-south” or in other words “in and out” of a data center. With modern applications such as big data modeling there is much more so called “east-west” traffic or in other words traffic that goes from one server (or virtual server) to another within the same data center. This fundamental shift in computing has led to new two tier data center fabric architectures such as leaf and spine (see **Figure 3**) where there are leaf switches which aggregate server traffic and switch it to other servers that are either accessible on the same leaf or must be passed to higher layer spine switches which are connected to every other leaf switch in the architecture. This architecture supports efficient, low latency communication between servers in a data center while also facilitating outbound “north-south” traffic.

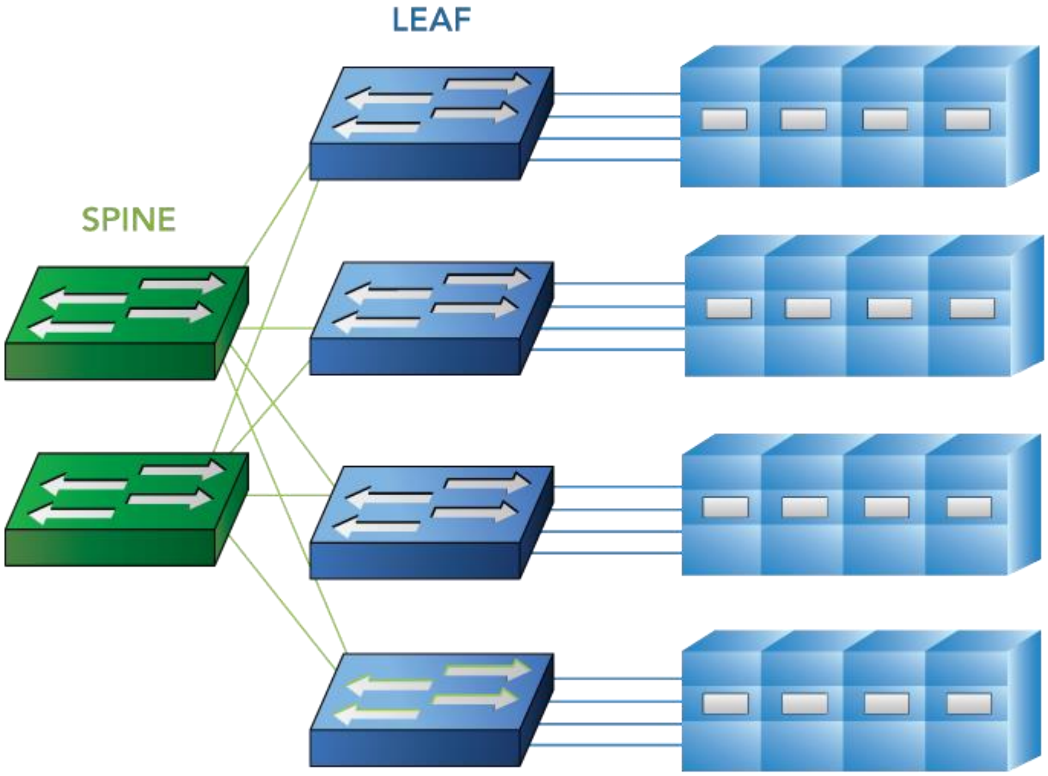

Figure 3: Leaf-Spine Architecture

In 2014 the first fixed configuration leaf switches (also called “top-of-rack” or TOR switches) with 100GbE uplinks emerged, paving the way for large scale, cost efficient 100GbE data center deployments in the future. Sophisticated deployments which use 10GbE connected servers for HPC or storage are the type of applications most likely to utilize these 100GbE uplinks in the short run. Clearly switch vendors such as Cisco, Juniper, Arista and others are setting the stage for 100GbE in the data center.

Not all 100GbE is Created Equal

As discussed, 100GbE products are rapidly emerging from many vendors with deployment in the telco carrier infrastructure and enterprise data center environments. However, customers must diligently evaluate new offerings to make sure they perform to specifications.

In the early phases of a product's lifecycle it is not uncommon for marketing hype to get ahead of technical reality. Vendors often have to make compromises because they have not resolved all the technical glitches inherent in nascent 100GbE products and are under revenue pressure to release products not ready for prime time. For example, a network security or monitoring application may perform well at 10GbE speeds but become non-deterministic at 100GbE and suffer performance degradation at line-rate. One way to mitigate these issues is to implement hardware acceleration to provide offload relief for the software application and associated host CPUs so they can perform processing tasks without being overwhelmed by the deluge of data.

Accolade 100 Gigabit Ethernet Acceleration Adapters

Accolade continued its tradition of technology leadership in lossless packet capture and acceleration with the introduction of the market's first dual port 100GbE adapter, the ANIC-200K in October, 2014 and the ANIC-200K4 (with CFP4 optics) in April, 2015.

The ANIC-200K4 captures data at 200 Gbps line-rate and buffers traffic in its on-board memory. This FPGA-based hardware adapter is easily integrated into network appliances such as test and measurement equipment that requires lossless data capture and traffic analysis at 100GbE speeds. For instance, testing and provisioning of 100GbE carrier Ethernet requires uncompromised packet capture at line rate, precise time stamping and latency measurement.

Accolade Hardware Adapters

Accolade's ANIC line of application acceleration adapters are fully PCIe compliant and seamlessly integrate into standard servers and network appliances offered by industry leaders such as Cisco, Dell, HP, Super Micro and many others. Accolade's OEM customers are leading network appliance vendors offering products for network monitoring/forensics, deep packet analysis, video stream monitoring, network test and measurement, high frequency trading (HFT) and high performance computing (HPC).

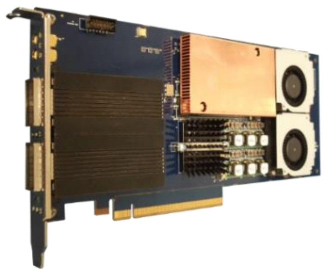

Figure 4: ANIC dual port 100GbE Adapter

ANIC adapters come in a variety of port configurations with speeds ranging from 1 to 100 Gbps. Each adapter performs lossless packet capture and a variety of advanced application acceleration functions including time stamping, packet filtering, flow classification and much more. Captured network traffic is buffered onboard the adapter in order to optimize the transfer capability of the PCIe bus. The network traffic is then intelligently steered directly to the appliance's (host) memory for processing by the software application. For multi-core CPU based appliances, the adapter intelligently distributes traffic to all available CPU cores thereby maximizing efficiency and performance.

Accolade provides its global OEM customers unparalleled hardware and software product integration support. A single, common and fully optimized software API across the entire ANIC product line enables seamless migration from one adapter to another ensuring quick time to market and investment protection.

Company Profile

Accolade is the technology leader in high performance, FPGA-based, lossless packet capture and application acceleration adapters. Accolade serves the global network appliance OEM market. Customers integrate the company's ANIC adapters in to their network appliances in order to gain advanced capabilities such as line rate packet capture, time stamping, packet filtering, and flow classification. Established in 2003, Accolade Technology is a privately held company based in Massachusetts with additional offices in Silicon Valley, California and Atlanta, Georgia.

Corporate Headquarters:
124 Grove Street, Suite 315
Franklin, MA 02038
T: 877.653.1261
F: 208.275.4679

Silicon Valley:
980 Mission Court,
Fremont, CA 94539
T: 877.793.3251

South East U.S. Regional:
2997 Cobb Parkway,
Atlanta, GA 30339
T: 877.897.4833

www.accoladetechnology.com