100G+ on a Standard Platform

Doug Sandy, Chief Technology Officer, and Todd Wynia, Vice President for Communication Products

October 2014

Smart devices and cloud services are driving dramatic increases in network bandwidth. Behind these services is a network infrastructure that needs to simultaneously deliver security, flexibility and performance. Security and flexibility require line rate deep packet inspection (DPI) and support for software defined networking (SDN) and network functions virtualization (NFV). Performance requires support for multiple 10Gbps, 40Gbps and 100Gbps interfaces. The systems that support this network infrastructure need to be scalable and upgradable so service providers can quickly respond to demand and deliver profitable services to satisfied customers.

This white paper discusses the bandwidth challenges being experienced by service providers and explains why data traffic will continue to grow. It goes on to discuss the benefits of using an AdvancedTCA® (or ATCA®) platform for telecom and other high availability applications and explains two approaches to increasing ATCA platform capacity.

The paper concludes with an overview of the Artesyn Centellis® platform and explains how it can be used to support up to 1.6Tbps aggregate data bandwidth today and up to 4.0Tbps in the future using the same chassis and backplane for applications such as deep packet inspection, security, packet classification and load balancing.

ARTESYN™
EMBEDDED TECHNOLOGIES

Bandwidth Challenges

Data usage across all parts of the network is growing rapidly. This growth is particularly strong in the mobile network where smart devices are driving exponential bandwidth growth. The Cisco VNI Mobile, 2014 forecast (see Figure 1) shows mobile data growing by a factor of 10 over the next five years giving a compound annual growth rate of 61%. This capacity growth rate translates to additional capacity demands in most areas of the network. As a result, operators are forced to either scale up their architectures by increasing the capacity of their equipment or scale out their architectures by adding more equipment, or both.

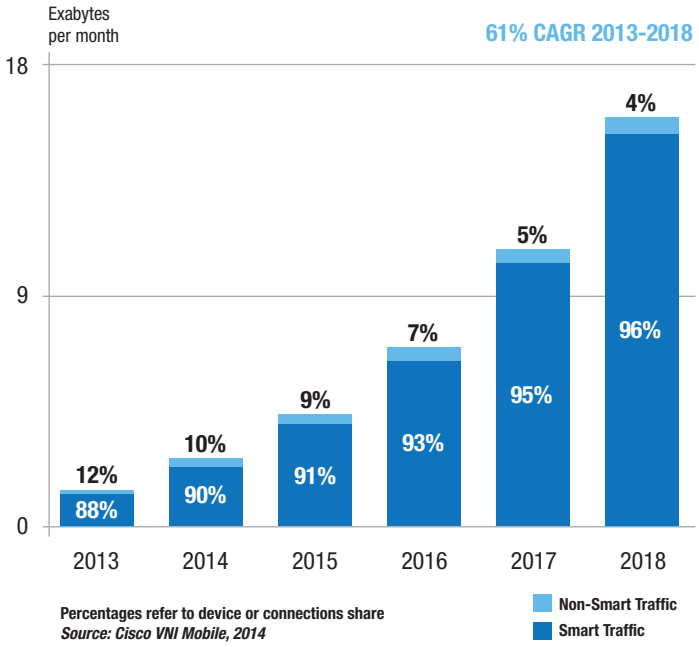

Figure 1: Growth of Mobile Data

The network and server infrastructure for mobile and fixed line services requires many line rate functions. These include security, deep packet inspection (DPI), packet classification and load balancing. Line rate DPI enables advanced policy management and gives service providers full control of network resources for individual users and user groups. As network capacity increases, these line rates must also scale upward from 10 and 40Gbps today to hundreds of gigabits per second in the near future.

Operators are looking for a platform architecture that will help address the need for scaling while also providing flexibility, 5 nines availability and an open, multivendor ecosystem. With its proven track record, having been deployed in most of the world's major networks, and a strong multi-vendor ecosystem, ATCA is a solid choice.

Platform Architecture

ATCA is a standardized platform for telecom and computing applications that has proved to be very flexible. Initially developed in 2002 ATCA is now used by many leading telecom equipment makers and is deployed in most of the world's major telecom networks. As an open architecture, ATCA has a broad multi-vendor ecosystem of blade and system options to choose from.

Distinct from most rackmount server architectures, ATCA platforms support up to sixteen slots for any mix of CPU, DSP, storage and I/O payload blades, together with up to four switch blades. Each blade can also have a smaller rear transition module (RTM) located behind the backplane to support additional storage or I/O capacity. This isolates the major I/O from the compute blades, easing cabling and maintenance. The most common ATCA platform configuration is a 14-18u shelf with 14 slots for mounting in a 19-inch rack.

ATCA platforms integrate dual power supplies, dual shelf management modules and multiple fan trays. The ATCA shelf manager controls all the modules and blades in the platform, supporting automatic switchover from active to backup switch blades, power supplies, cooling trays and shelf managers as needed. This approach allows ATCA systems to meet the requirements of Network Equipment Building Standard (NEBS) for telecom and other high availability applications.

ATCA platforms have a passive backplane that supports Base and Fabric Interfaces. The Base Interface is a pair of 1 GbE channels that are connected from each payload blade to a pair of centralized switch/hub slots. This Dual Star configuration supports active and standby switch blades. In telecom systems the Base Interface is typically used for control plane functions.

The Fabric Interface is the high-speed interconnect within the ATCA platform and is typically used for data plane functions. The Fabric Interface can be implemented as a mesh or a star configuration. Several star configurations are shown in Figure 2.

Most current ATCA platforms support the Dual Star Fabric Interface configuration. In a 14 slot shelf this supports 12 payload slots and 1+1 redundancy with one active and one standby switch blade. The main system Ingress/Egress ports are usually connected directly into the active switch.

The 2+2 Dual-Dual Star doubles the system switching and Ingress/Egress bandwidth with two active switch blades and two standby switch blades. The number of payload blades in this configuration is reduced by two to accommodate the additional switch blades.

A more efficient implementation with four switch blades can be achieved by using Artesyn's 3+1 QuadStar™ configuration. This configuration has three active switch blades supporting three times the bandwidth of a standard Dual Star system. The fourth switch blade is ready on standby in case one of the three active blades fails.

Each Fabric Interface channel consists of four bidirectional differential pairs (lanes) combined to provide 10GbE or 40GbE bandwidth. To achieve the best results 40Gbps ATCA platforms need to have 10Gbps KR lanes and IEEE40Gbase-KR4 switching.

All ATCA platforms support backwards compatibility to 10GbE and 1GbE. For future compatibility it is critical that the ATCA chassis has a backplane that can support KR level performance. For more information on ATCA see the Artesyn white paper, ATCA Yesterday, Today and Tomorrow.

Scaling ATCA to 4Tbps

An ATCA system with 40GBaseKR4 switching and Dual Star Fabric Interface will support 40Gbps to each payload blade. In a 14 slot system this provides a total aggregate bandwidth of 480Gbps.

ATCA can be scaled to support up to 4Tbps aggregate bandwidth. The first step is a QuadStar™ implementation with 40G fabric interfaces. This yields a non-redundant aggregate bandwidth of 1.6Tbps. To scale further, 100G fabric interfaces are required. Implemented in a QuadStar™ topology, 100G fabric interfaces yield 4.0Tbps aggregate bandwidth.

In the same way that the ATCA fabric interface has scaled from 1GbE to 10GbE and now 40GbE, future switches and payload blades will support 100GbE fabric interfaces with 25Gbps on each lane. This approach requires silicon switch devices that support 12-16 100GbE connections, a new connector and agreed standards for implementing 25Gbps lanes on ATCA.

The 3+1 QuadStar Fabric configuration with 40GBaseKR4 switching supports 120Gbps per blade or 160Gbps with four active switch blades and no redundancy, equating to 1.6Tbps aggregate bandwidth. This configuration can be implemented using current technology backplanes, switch devices and payload blades with support for four fabric interfaces.

Working with other industry leaders, Artesyn has pioneered the technology to enable 100Gbps backplane links within ATCA. This means that Artesyn's chassis and backplane supports 300Gbps interfaces with redundancy to each payload blade. In a 14 slot system this provides an aggregate data bandwidth of 3Tbps or 4Tbps with no redundancy.

Artesyn, the industry leader in deployed AdvancedTCA, has pioneered the availability of multiple 100G fabric interfaces in ATCA with its flagship Centellis® platform.

Centellis 8000 Series Platform

The Artesyn Centellis 8000 series is the industry's first high availability telecom platform for application acceleration and advanced networking that incorporates 100Gbps QuadStar™ backplane technology. (see Figure 3). The platform can be deployed in a wide range of applications including security, DPI, packet classification and load balancing in both physical networks and virtual networks supporting SDN and NFV. Application-ready configurations for several of these applications shorten time-to-market.

The Artesyn Centellis 8000 series is available with AC or DC power and cooling to support up to 600 watts/slot. Pre-certification of DC power configurations for NEBS Level 3 saves on testing costs and further shortens time-to-market. The 18u shelf has 14 slots for ATCA switch and payload blades.

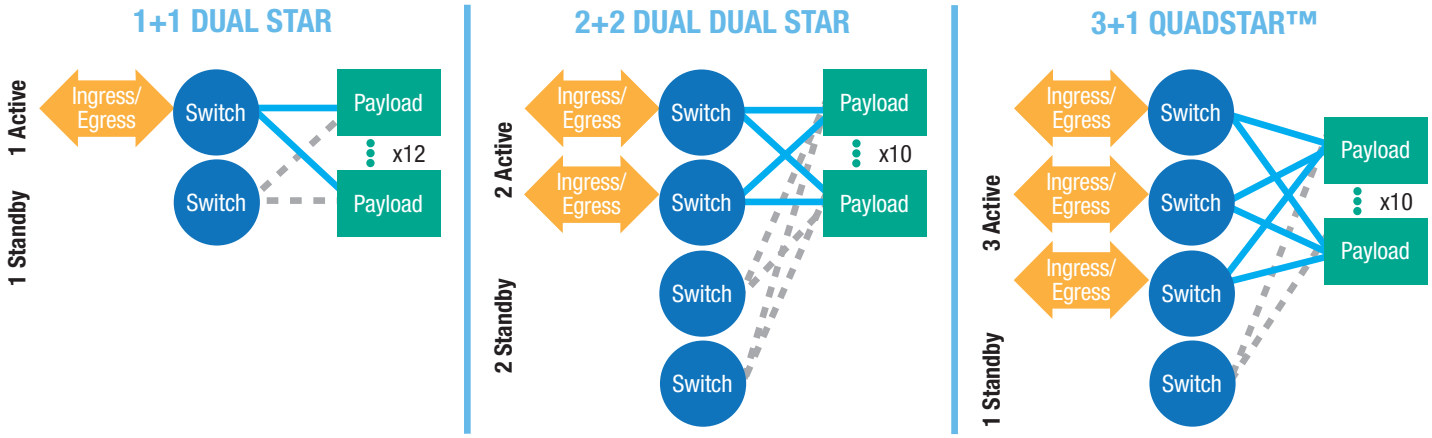

Figure 2: ATCA Fabric Interface Star Configurations

Figure 3: Artesyn Centellis 8000 series telecom platform

The Centellis 8000 series platform supports multiple backplane, switch and payload blade configurations. The performance of five configurations is summarized in Figure 4.

Config 1 shows a 1+1 dual star configuration with up to 12 payload blades and two 40Gbps switch blades, one active and one standby. Each payload blade has a 40Gbps link to the active switch blade supporting an aggregate bandwidth of 480Gbps for the whole platform.

Config 2 is a 2+2 Dual Dual system with two active 40Gbps switch blades and two standby switch blades. This configuration will support up to 10 payload blades with 80Gbps from each payload blade split across the two active switch blades. Aggregate bandwidth for this platform configuration is 800Gbps.

Config 3 is a 3+1 Quadstar™ configuration with three active 40Gbps switch blades and a single standby switch blade. In this configuration the payload blades have four 40Gbps links, three to the active switch blades and one to the standby switch blade. The bandwidth per payload blade is 120Gbps giving an aggregate bandwidth of 1.2Tbps for the whole platform.

Config 4 is a configuration with no redundancy and all four 40Gbps switch blades are active with four 40Gbps links to each payload blade. The bandwidth per payload blade is 160Gbps giving an aggregate bandwidth of 1.6Tbps for the whole platform.

Config 5 is a 3+1 Quadstar™ configuration with 100Gbps switch blades and payload blades with 100Gbps fabric interfaces. The bandwidth per blade is 300Gbps and the aggregate bandwidth is 3Tbps.

Config 6 is a configuration with no redundancy and all four 100Gbps switch blades are active with four 100Gbps fabric interfaces to each payload blade. The bandwidth per blade is 400Gbps and the aggregate bandwidth is 4Tbps.

All six configurations can be supported by the same Centellis 8800 series platform with QuadStar™ backplane. This provides significant scalability and future proofing. Platform bandwidth with redundancy can be increased from 480Gbps to 3Tbps without changing the chassis or backplane.

Application-ready configurations supporting Dual Star, Dual Dual Star and QuadStar™ redundancy with 40G fabric interfaces are available today. Configurations supporting 100Gbps switching and payload blades will be available in the future using the same chassis and backplane.

Conclusions

ATCA is a flexible platform for telecom and other high availability applications. The continued development of more advanced backplanes and higher speed switching, together with power and cooling up to 600W per slot, is increasing ATCA system capacity by a factor greater than 10, more than matching the growth in network traffic.

The Artesyn Centellis 8000 series telecom platform is already supporting system bandwidths up to 1.6Tbps with QuadStar™ configurations. By using new switch and payload blades the system capacity will grow to 3Tbps with 3+1 QuadStar redundancy and 4Tbps with four active switch blades.

For more information on Artesyn's ATCA solutions, please contact your local Embedded Computing sales office.

	Config 1	Config 2	Config 3	Config 4	Config 5	Config 6
Fabric Technology	40G	40G	40G	40G	100G	100G
Redundancy	1+1 Dual Star	2+2 Dual Dual Star	3+1 QuadStar™	None	3+1 QuadStar™	None
Payloads per shelf	12	10	10	10	10	10
Payload slot	40G	80G	120G	160G	300G	400G
Bandwidth available for each payload slot	480Gbps	800Gbps	1.2Tbps	1.6Tbps	3.0Tbps	4.0Tbps
Availability	Today	Today	Today	Today	Future	Future

Figure 4: Artesyn Centellis 8000 series Performance Summary

100G+ on a Standard Platform white paper R1D1