REPORT REPRINT

Arista launches first switches based on XPliant 'SDN' packet forwarding chips

PETER CHRISTY

13 DEC 2016

The company has announced the 7160 family of compact switches, with three models varying in port configuration, all to be available in Q1 2017. Predictably, the switches feature greater capacity (6.4Tbps, 1.28Bpps) and faster ports.

451 Research®

 |

Arista has announced the 7160 family of compact switches, with three models varying in port configuration, all to be available in Q1 2017. Predictably, the switches feature greater capacity (6.4Tbps, 1.28Bpps) and faster ports. The switches are based on the XPliant XP80 from Cavium (one of the new generation of more programmable devices that have been introduced as part of the SDN movement), new software features that leverage the chip architecture (specifically AlgoMatch for greater ACL capacity) and lots of nuanced details that highlight the non-obvious cost economics of optical networking.

THE 451 TAKE

The 7160 switches are unexceptional in the sense that we have grown to expect timely product-line upgrades from Arista that incorporate faster and larger switch parts, and that leverage Arista's software agility to add new features. The 7160 is notable because it is the first to incorporate one of the 'SDN' (highly programmable) chips (in this case, the XPliant chip from Cavium), and because Arista's exploitation of the new part is more nuanced than the earlier vendor positioning that emphasized the ability to incorporate new protocols in the switch chip (a value proposition that is still speculative). Arista is a good early indicator of what's happening in the rapidly evolving world of cloud datacenter networking, given that the hyperscale providers themselves tend to be mum. In this case, Arista described the continuing growth in datacenter bandwidth and the attractiveness of 25Gbps and 50Gbps links over 40Gbps because of the cost economics of the electro/optical transceivers.

CONTEXT

Arista is the most successful network equipment vendor of the last decade, having built a $1bn switch business in a difficult market. For the last two years, it has been under legal attack by Cisco for alleged misuse of Cisco IP. When Arista was founded as Arasta, it was a software company. It then transformed into an equipment vendor to capture a larger profit stream to drive ongoing investment.

A key strategy decision was to forego the custom silicon development done by the leading vendors, and depend instead on the use of merchant silicon, complemented by the advanced software architecture (EOS) and the system design elegance we expect from Andy Bechtolsheim, the funder and early product leader.

Arista says that EOS, with its unique, componentized software architecture, is a key competitive differentiation because it enables the company to develop new switch features and software releases considerably more quickly than its competitors, which in the last 18 months have included new multi-switch distributed system features that enable simplified network management and instrumentation.

PRODUCTS

The announced new switches were the 7160-32CQ with 32 QSFP100 (100Gbps) ports, the 7160-48YC6 with 48 25G ports and 6 100G ports, and the 7160-48TC6 with 48 10G-T and 6 100G ports. All do wire-speed L2 and L3 with VXLAN, VXLAN bridging and routing, 128-way ECMP with IPv4 and v6 routing, 2-3 usec latency with up to 6.4Tbps and 1.2Bpps throughput.

All run the same single-image Arista EOS software, and all incorporate new software/hardware functionality called AlgoMatch (that depends on the new chip architecture) to provide a larger ACL capacity for richer network policy support with 2-6x the capacity of traditional methods, and lower power consumption than traditional TCAM implementations.

Arista says the new chips enable the most flexible implementation of dynamic buffering, yet have the ability to allocate packet buffer memory on a per-line basis based on dynamic load, additional flexibility in partitioning switch resources for L2 or L3 as needed for specific customer traffic, a programmable pipeline, and support of the full IEEE 25GbE 802.3, which enables rack-scale network cost performance cable tradeoffs.

MARKETS

Arista has emphasized 'cloud networking' from the beginning, wisely choosing what turned out to be the growth segment in switching, and one in which full Cisco protocol support was not required. However, over time, Arista (and all other switch vendors) are being squeezed out of the volume ToR hyperscale market by white box or internally developed switches.

Arista continues to sell products into that market, and is a leading vendor in lower-tier cloud switching markets. Close ties to the hyperscale providers are valuable, because of the speed with which they incorporate and evolve new technology and because of their relentless focus on cost and performance optimization. Whereas 15 years ago, the leading network vendors engineered the most advanced networks, in the last 10 years, that role has increasingly been done by the hyperscale systems operators.

TECHNOLOGY

The 7160 is the fourth Arista 'architecture' (defined by the merchant switch family used internally). The first was Fulcrum-FM based (Intel), and then there have been two based on Broadcom (DNX and XGS). This new architecture is based on the XPliant parts (Cavium acquired XPliant in 2014).

XPliant was the first 'SDN' switch part startup with a design that better enabled OpenFlow switch programming. When OpenFlow was introduced as the centerpiece of the Open Networking Foundation, the idea was to revolutionize network device implementation, replacing the switch internal software with an external controller running OpenFlow – intended to be a generic 'machine language' for programming switch parts. The existing merchant parts (e.g., Fulcrum, Broadcom) were simpler (in terms of logic complexity), and better understood as the least common denominator of the various network software stacks in the market.

The XPliant parts implemented the same logic functions as existing parts but added flexible interconnection logic between those major functions that didn't exist in the existing merchant parts. Although this added considerable complexity to the part (more transistors), XPliant (and other advanced architecture chip vendors) argue that it doesn't increase the chip size or cost because that was defined by the I/O requirements, and there was 'spare' space to add this new logic without increasing the chip size.

The potential value of this increased programmability was clear (and not available) when networking was developed for virtualized servers. Adding virtualization added new network endpoints – the virtual machines running on a virtualized server, each of which thought it was directly connected to the network – and that in turn required adding a new level of network addressing and new addressing protocols (VXLAN).

Server virtualization occurred rapidly because existing software workloads continued to run unchanged, creating the need for virtual networking far in advance of the normal time frame in which new technology is added to networking chips. As a result, virtual networking was initially done by a new layer of software switches added at the hypervisor (vSwitch), with the physical network not materially contributing. When new chips with VXLAN awareness (the ability to act on the new packet headers) became available, switches could play a more active role in VXLAN networking.

VXLAN is the marketing argument given by XPliant (and now repeated by Arista) as to the intrinsic value of this kind of flexible architecture (although the commercial value in the Arista switches is the ACL and L2/3 flexibility). If new protocols are developed and implemented as rapidly as was the case with VXLAN, built-in chip architectural flexibility will clearly be valuable. If the protocol changes with a longer time frame, it will often appear in 'conventional' silicon in an adequately timely fashion. Only time will tell whether VXLAN was a unique event, or a harbinger of more agile networking to come.

Arista says the greatest commercial value from the new chips (enabled by new Arista software) is in providing significantly expanded ACL capacity (2-6x the competition according to Arista), while saving power. Arista does this via a new hardware/software function it calls AlgoMatch, that looks up ACLs in a larger SRAM store, rather than using more expensive and power-hungry TCAM associative memory. It uses a hashing-like technique for the lookup (saving power because less of the memory has to be searched).

Arista argues that cloud computing requires a larger number of ACLs and more dynamic ACL management, both of which AlgoMatch provides. Arista also says the new hardware enables a more flexible adjustment of switch resources between L2 and L3 actions, and more flexibility in the assignment of packet buffering based on dynamic need.

Finally, Arista says that it anticipated a rapid evolution (by historical standards) from 10G to 25G as datacenter communications continue to grow rapidly, driven by the continuing improvement in CPU power, and by the increasing 'East/West' traffic of modern applications.

In the datacenter, 25G is surprisingly cost attractive because copper (twinax) server to ToR links still work for 25G (although according to Arista, they often need the forward error-correcting features of the full 25GE standard). Electrically, 40G is implemented in four 10G lanes, whereas 25G is a single signal, so the high cost of the electro/optical transceivers makes 50G (two 25G lanes) attractive, compared to 40G.

COMPETITION

Arista's direct competitor is Cisco, the switch industry leader, and now it is clearly in Cisco's competitive crosshairs. Arista also competes with HP, Dell, Juniper, Brocade and others in the switch business, but in recent years, none of these have kept up with Arista's growth. In cloud computing (the most rapidly growing part of the switch business), Arista competes with in-house implementation in the case of the largest services (it has been estimated that the top seven services represent more than 50% of the total market).

SWOT ANALYSIS

STRENGTHS

Arista has demonstrated the ability to deliver a continuing set of innovative products, and grow revenue and market share in a difficult market.

WEAKNESSES

As Arista gets larger, is it more subject to the difficulties of the market, and more directly subject to competition with Cisco.

OPPORTUNITIES

Cloud datacenter networking continues to evolve rapidly in terms of line rates and switch throughput.

THREATS

One of the greatest threats to the switch business as a whole is the continuing growth of the hyperscale cloud providers, and their increasing ability and desire to build their own networking systems.