Big iron, fast photons and SDN: the hard, bright and soft sides of Arista's strategy

Analyst: Eric Hanselman Peter Christy

2 Jul, 2013

It's been a busy spring for Arista Networks, with expanded software functionality, upgraded switch modules and a Best of Interop/Grand Prize win for its 7500 Series modular switches. The new software capabilities include integration of its Arista EOS switch operating system with OpenStack (enabling fully automatic provisioning); a contribution to the OpenStack Neutron networking project (formerly OpenStack Quantum); OpenFlow support with a controller-less option; and a significant enhancement of its eAPI interfaces for EOS. The hardware upgrades encompass big increases in capacity and port density for its 7500 Series switches, as well as hardware VXLAN support and an approach to 100Gb Ethernet interfaces that breaks new ground in terms of price/performance and flexibility. All of these enhancements are delivered with what Arista pitches as a stress-free migration path for existing customers.

The 451 Take

While Arista's original product designs could have claimed to do SDN 'before SDN was cool,' the vendor's latest announcements extend integration into the major SDN camps of OpenStack and OpenFlow, and also increase programmability. These enhancements can offer value in both simplified datacenter cores and complex multi-tenant networks. Arista is also building on its high-performance legacy, leveraging its relatively modern, clean software design and merchant-silicon-based hardware to add a level of capacity in its large-chassis switch products that should challenge rivals. A simple upgrade path will please Arista customers, but the true innovation value may lie in a gamble on 10/40/100Gb Ethernet optics, whose lower cost could really pay off.

Copyright 2013 - The 451 Group

Context

Arista was founded in 2004 and funded by successful entrepreneurs Andy Bechtolsheim and David Cheriton. From the beginning, the vendor placed emphasis on the software architecture of network switches, choosing to use commercially available ASICs rather than developing its own. At the beginning of 2013, Arista had grown to more than 500 employees and 1,800 customers. The company claims it is profitable, and many expect an IPO in the not-too-distant future.

Arista has focused on a sequence of markets where its systems engineering and software platforms can be leveraged to succeed against much-larger networking industry incumbents – including education, high-performance computing, high-speed trading, network performance and, most recently, network monitoring systems.

Technology

Arista's recent software announcements fit well within 451 Research's definition of software-defined networking, which embraces a broad set of approaches to improve the programmability of networks, often toward the goal of improved automation of network management (not just the use of OpenFlow.) With these moves, Arista enhances its support for the leading standards-based approaches, OpenFlow and OpenStack, to complement its advanced integration with VMware (it has aggressively introduced features to support VXLAN virtual networking) as well as leveraging its internal software architecture to provide API support for a diverse set of industry partners, including Riverbed, Palo Alto Networks, Aruba Networks and Splunk.

The first key technology here is EOS, Arista's relatively clean and modern (compared with the larger incumbent vendors) internal software. EOS is based on standard Linux, and the higher-level management functions within a switch can be implemented as Linux applications, both simplifying their construction and making it possible to extend an Arista system with mechanisms familiar to server admins.

The design of EOS has value as the basis for providing SDN capabilities, as the extensions to the EOS eAPIs demonstrate. Modern switches provide rich management capabilities via CLI commands – a network administrator connects to a specific device and performs management functions via text commands (reminiscent of how PCs were used during the DOS era). Cisco is creating programmatic equivalents to a subset of its own suite of CLIs (onePK) as an important element of its SDN strategy, pointing out that many 'SDN' applications can be built using these interfaces.

Arista's eAPI offering is analogous to Cisco's onePK, but is more comprehensive. Cisco has many products within its current portfolio, and a diversity of specific CLIs and functions as well. OnePK builds a programmatic interface to a common subset of those commands with a single set of functions across all these products – limited intrinsically by the differences between the products, however. Arista, which started much later and views software as a key asset, has a simpler task here because a single EOS image is shared by all Arista products, and because EOS has included API interfaces from the beginning. Arista enhanced EOS so that CLIs provide a JSON-based programmatic equivalent. It claims there is over 90% coverage of CLIs in the initial eAPI release, with full coverage to follow shortly. JSON offers an advantage over XML-based APIs, since it eliminates the risk of parsing errors in malformed XML structures.

Neither Arista nor Cisco claims that these adaptations of CLIs are a complete alternative to what OpenFlow might eventually provide. Instead, they say that such commands are an alternative approach to many of the problems that OpenFlow is said to address – an approach that leverages existing switch functionality rather than requiring that it be re-implemented within the OpenFlow architecture.

OpenFlow and OpenStack additions

Arista is also embracing OpenFlow. It offers version 1.0 support on the 7050 chassis and has tested interoperability with controllers from Big Switch, NEC and others. Most switch vendors (now including Arista) offer one or more switches that have been adapted to enable OpenFlow control. The use of switches in a pure OpenFlow mode (where the legacy autonomous mode has been disabled) has been largely limited to research purposes, outside of hyperscale users such as Google. A larger number of applications have been able to use the switch in a hybrid mode of operation where the existing control-plane protocols (e.g., routing protocols) manage most of the packet-forwarding rules, while OpenFlow is used to create specific routes or manage certain ports.

Arista has also introduced a new twist on OpenFlow use with its recent announcements. Arista's DirectFlow is an interesting angle for the specific and limited use of OpenFlow-injected rules. DirectFlow enables rules to be added to the packet-forwarding tables in an Arista switch (that are still managed by the legacy autonomous switch operation), but rather than worrying about how to integrate these two rule sets, DirectFlow simply invokes the OpenFlow rules first. If an OpenFlow-inserted rule matches an incoming packet, then it specifies how the packet will be forwarded. If none of the OpenFlow rules trigger the rule set, then the legacy switch functions are used. Arista calls DirectFlow 'controller-less OpenFlow,' since the rules can be configured directly on the switch, rather than requiring an OF controller, which itself can offer considerable cost savings.

While OpenFlow offers low-level control, Arista is also adding integration with the OpenStack cloud orchestration framework. It's contributing software to the OpenStack Neutron networking project that translates network allocation requests into network configurations. Plug-ins for Neutron are available for a number of vendors and purposes, and Arista is contributing code to the open source project to allow both virtual and physical networking devices to be controlled at the same time. This simplifies orchestration operations that would previously have to use different plug-ins for different underlying networks. It expects this capability to be included in the 'Havana' release of OpenStack this fall.

Optical interfaces

Arista has taken an innovative approach to increasing port density and decreasing cost for its fiber-optic interfaces. The shift to optical interconnects has fractured interface choices, with differences in speeds, as well as fiber connectors and diameters. The response of equipment vendors has been the use of SFP (small form-factor pluggable) modules that allow simple swapping to match the requirements of the installation. While this increases flexibility, it also increases cost and creates a procurement headache in ensuring that the right types of modules are on hand.

To reduce cost and, hopefully, speed adoption of 100Gb Ethernet, Arista has taken a rather bold step by using integrated modules with direct fiber connections, rather than SFPs. The initial release, in the 7500E, uses multimode fiber – the type used for shorter-distance connections typical of datacenters. There is an intimation that single-mode fiber – the type used for kilometer-scale connections – might be available in later products, but is not part of this release.

This decision is a bet that datacenter networks will be the major volume opportunity for 100Gb Ethernet, and that an attractive price point will trump any concerns about flexibility. This seems like a safe wager. To stack the deck in its favor, Arista allows each interface to operate in 10Gb, 40Gb or 100Gb modes. Both 40Gb and 100Gb Ethernet have versions that use multiple 10Gb lanes, making this option possible. This flexibility will smooth the transition to 100Gb by allowing customers to use the same switch blade for current 10Gb requirements and then shift ports to 40Gb or 100Gb as needed.

VXLAN beyond virtualization

There are a number of network tunneling protocols that have grown out of the various virtualization approaches. VXLAN has been used by VMware to extend virtual Layer 2 networks across routed Layer 3 networks. Arista has added the ability to not only act as a gateway for bridging between VXLAN tunnels and configured networks or VLANs, but to also use VXLAN to extend Layer 2

networks between physical switches.

The advantage of using VXLAN is that it offers over 16 million virtual network identities, compared with the 4,096 available in the traditional VLAN approach. This enables the support of very large virtual infrastructures, or very large numbers of tenants for service providers. This also circumvents the issue of creating fabric architectures, such as Cisco's FabricPath, Juniper's Qfabric and Brocade's VCS. Users can build Layer 3 network architectures that hyperscale operators have been using to scale efficiently, and then create Layer 2 virtual overlays as needed.

One of the scaling issues with encrypted tunnels like VXLAN is the amount of computational work required by the encryption and decryption at tunnel endpoints. In a hypervisor-managed tunnel, that work is done by the server processors. Arista offloads that work to the switch hardware for the tunnels that it manages, allowing greater scale.

Products

Arista is rolling all of these technological enhancements into updated blades for its 7500 Series chassis switch with greater performance. The 7500E series bulks up forwarding performance to a claimed 14 billion packet-per-second rate and increases the backplane capacity to 30Tbps. To manage bursty traffic profiles, the system can be configured with as much as 144GB of packet buffer memory.

Four different line cards are available with different port configurations in this higher-performance 7500 Series. The 7500E-48S offers 48 ports of 10Gb Ethernet using traditional SFP+ optical modules. The 7500E-36Q has 36 QSFP+ ports that can be used as a mix of 40Gb Ethernet (up to 36 ports), or any QSFP+ can be broken into four 10Gb interfaces, for a total of 144 ports in an all-10Gb configuration. The 7500E-12CM and 7500E-72S models offer 100Gb ports, using 12-strand MPO connectors that connect directly to the embedded optical modules. The former has 12 MPO ports that can be all 100Gb or broken out into 10Gb or 40Gb links, while the latter has two MPO connectors and 48 SFP+ modules. The SFP+ ports are 10Gb only, while the MPO ports can also be broken out into 10Gb and 40Gb, as well as 100Gb.

These new cards make upgrades simpler, since they can be added to an existing 7500 chassis. Total performance will be limited with older cards in the mix, but customers can upgrade incrementally to spread out the costs. There is no change to the existing chassis or power supplies, and the line cards, supervisory module and chassis fabric can be upgraded independently.

Per-port costs for the new systems when fully loaded are said to be around $10,000 for a 100Gb

Ethernet port, including optics. That's a fraction of what has been the typical cost of ports using separate optics in the CFP form factor. Costs for 40Gb and 10Gb ports in similar configurations are around $2,200 and $550, respectively. These price points have the potential to markedly shift the economics of network design.

Competition

Arista's primary competition in datacenter switching is market leader Cisco, followed by HP, Juniper and others. HP claims the lead in shipping OpenFlow-enabled ports (citing over 15 million) but, to date, the capability hasn't reached its datacenter-class products. It's also uncertain how many of these potential HP ports are actually in production using OpenFlow. Arista's release joins capabilities that Brocade offers in its MLXe routers, Juniper offers in its MX and EX products, and the proof-of-concept support that Cisco is offering in Nexus and ASR.

Raw networking performance continues to be Arista's strength. The capacity of the 7500E in port density and traffic handling puts it in rarified air. Avaya and Alcatel Lucent have long histories with network operators and service providers, but will struggle to match the new price/performance levels that Arista is setting. The bevy of ODM and OpenFlow network hardware vendors – including Pica8, Pluribus, NoviFlow and Cumulus – may introduce some pricing pressure, but Arista's capabilities should maintain solid differentiation from any of these products. Intel's network platform efforts, meanwhile, will likely only enhance Arista's position, given that it's a customer of Intel's Fulcrum merchant silicon.

There is additional competitive pressure (albeit early on) from virtual network offerings. Virtual networks are reducing port count requirements and converging larger volumes of traffic on a smaller number of higher-speed ports. It looks like Arista, which has focused on higher performance, aligns well with this shift. Having greater density in 10Gb, 40Gb and 100Gb ports matches the increases in server port speeds and greater network backbone scale.

With the enhancements to its eAPI offerings, Arista takes a step up in offering pragmatic SDN alternatives to the use of OpenFlow. This has become an area of interest as users gain more experience with the realities of OpenFlow deployments.

Altogether, this raft of new capabilities keeps Arista solidly placed in the datacenter and networking core. The combination appears well aligned with a datacenter focus, and the optical interfaces could lead the way to greater success.

SWOT Analysis

Strengths

Arista continues to make effective use of merchant silicon and thereby avoid the delays and risk of custom ASICs, while demonstrating value in its software-focused product design and innovative system designs.

Weaknesses

Although rapidly growing and achieving success as measured against many of the small network vendors, Arista is much smaller than market leader Cisco and has yet to demonstrate significant impact in the markets of greatest interest to Cisco.

Opportunities

The combination of VXLAN and OpenStack support, coupled with enhanced programmability, open the door for greater integration in virtualization and cloud environments and deeper partnerships with platform vendors. These are areas that, once established, will be hard to displace.

Threats

Cisco will continue to battle with performance improvements, and ODM networking companies will try to erode margins from below. Arista must demonstrate that its approach maintains differentiating value in ever more virtual networking environments.

Reproduced by permission of The 451 Group; © 2013. This report was originally published within 451 Research's Market Insight Service. For additional information on 451 Research or to apply for trial access, go to: www.451research.com