ARISTA

Featuring

FORRESTER®

Arista Q&A Document with Andre Kindness

on Data-Driven Networking

Arista recently conducted a video interview with guest speaker Andre Kindness. As a follow-up to the event, we had the opportunity to ask Andre more questions, about the future of the data-driven enterprise.

1. **What are the top three networking/infrastructure problems that you believe AI or machine learning (ML) can solve for enterprise customers?**

   Artificial intelligence and machine learning can bring immense value to the networking organization. While the market is relatively new, Forrester believes organizations can, today, extract from the technologies by using them to:

   a. Find problems. Finding the exact root cause of network outages and performance issues can bog organizations down for long periods of time. Networking professionals have the awful and cumbersome task of trying to figure out the issues and location of data, aggregate it, and set up tables to correlate events while dealing with decaying information. Real-time data lakes allow networking organizations to combine all sorts of data, including machine data (such as push and pull data along with logs); traditional IT service management (ITSM) data (tickets, change controls, and asset information); and outputs from network controllers or orchestration software. Contextualize the data and correlate automatically within the platform by using ML algorithms.

   b. Proactively suggest solutions. When it came to fixing, upgrading, or updating networks, an imaginary reactive wall existed between the client and vendors. Customers would call in with an issue, and a vendor's support team would work on resolving it. If the solution was a software upgrade to a vendor bug, other clients would only find out if they had the issue or read through notices. AI/ML transforms a reactive, transactional relationship into an active business partnership. By combining a vendor's AI solution with all the client data feeds, it allows the system to find errors, develop corrections, and push out fixes before clients stumble across them.

   c. Improve efficiency. Previously, network capacity management tools used analytics that focused exclusively on historic infrastructure utilization. With modern workloads, however, whether private cloud, public cloud, or hybrid, the past isn't always a reliable guide to the future; AI/ML solutions allow networking organizations to build whole systems that can scale up and down even when demand appears to have no discernable pattern to the human eye; inject new scenarios to see what happens before moving into production; and highlight missing or misconfigured settings that are hampering the network's performance.

2. **Do you believe AI and/or ML can autonomously administer a network? Why or why not?**

   AI and ML solutions provide many benefits, but running an entire business network autonomously isn't one of them. It's hard to find any large and complicated process that is 100% autonomous. Most systems that run an autonomous process are usually a subsystem that is either well defined, has a small amount of inputs, or is relatively isolated and has a small domain. None of these aspects reflect elements of today's networks that are supporting digital businesses. Moving forward, networks will only become large, more dispersed, and more virtual across personal and business

infrastructures. New inputs and variables will continuously be added. Instead of autonomous networks, networking professionals will guide the autonomous networking subsystems through these network changes and additions, reflecting an augmented network. A similar concept is fly by wire on modern aircraft that is combined with closed-loop systems: The aircraft engineers acknowledge that they can't create autonomous systems that account for every aspect an aircraft will encounter in its lifecycle and designs autonomous elements; this augmented aircraft enables the aircraft to optimize inputs made by the pilot and to prevent the pilot from operating outside of the aircraft's safe performance envelope.

3. **What do you believe is the right architecture for core/edge routers? Scale out or scale up? Why?**

   Scale-up architectures provide a short-term solution for the lack of a missing key component. A scale-up solution provides a single location that provides a single interface and location where all functions exist. This approach has many limitations, such as brittleness, high latency, and rigidness. For example, networking organizations have struggled to upgrade firmware and hardware with large routing chassis — either because it's hard to find a time to take them offline or because they can't add new line cards with existing ones. Distributed processes and networking architecture scale out to be more effective while overcoming the changes of scale-up designs. Besides the flexibility of making incremental changes to match the business without taking the network down, scale-out systems provide lower-latency communications and linearly match performance requirements while maintaining a linear cost equation — though a scale-out approach requires a solid set of policies, management and monitoring tools, and processes.

4. **Watching the evolution of SD-WAN technologies, do you feel there is a next generation coming? If so, what will its hallmarks or keystones be?**

   SD-WAN had its day in the sun and is now evolving to next generation routing WANs. Why? The concept was just a step in a bigger movement. The rise of SD-WAN highlighted a few trends underserved by the past WAN market: the importance of a highly reliable WAN to serve the digitalization happening inside hospitals, manufacturing plants, and retail stores, to name a few; the shift of focus to individual applications and their locations from traditional connectivity metrics; and being able to help create easier methods of managing the rise of businesswide networking fabrics. Much of the SD-WAN offerings didn't acknowledge the need for interoperability, advanced security services, sophisticated monitoring, or other network services such as WAN optimization or routing. The industry has started to evolve to weave together more network services and embed more security elements into the network fabric — Forrester calls this Zero Trust Edge. Yet the industry has been very quiet on the interoperability front. With hybrid and multicloud connecting the edge together, a single solution won't be able to accommodate all the facets of this complex WAN fabric. Forrester believes that organizations will require interoperability and openness from the vendors providing WAN fabric solutions, which will force the industry to create an open, secure set of WAN solutions.

5. **Is there value in normalizing networks to common design patterns and consistent NetworkOS/OAM&P choices for implementation of these design patterns in an infrastructure/network-as-code operational modality?**

   The one-offs are automation and efficiency killers. This is why standardization is one of the top elements that networking organizations should focus on. While businesswide networks cross many domains and require many different types of solutions, it is important to standardize on certain elements. This doesn't mean standardize on a vendor. A single networking vendor might offer different solutions, such as operating systems, across the data center, campus, cloud, or remote office. The reality is those groupings don't exist anymore. The businesswide network interweaves applications, data, and microservices across data centers, cloud platforms, and edge compute locations. I&O pros, orchestration and automation software, and monitoring systems all require consistent networking features and interfaces. This means that a consistent operating system within cloud, virtual, and hardware solutions has to be standardized. The Forrester Wave™: Open, Programmable Switches For A Businesswide SDN, Q3 2020 focuses on evaluating operating systems, not hardware, that enable a consistent fabric across the entire business network.