Remember when…it was sprawl or nothing?

Lori MacVittie, 2009-23-03

Ah, those were the days, weren't they? When improving the security, reliability, and performance of applications over the LAN, over the WAN, and over the Internet meant you had to deploy many different solutions, each one standing on their own in the data center. When you had to learn how to configure and manage as many devices as you have fingers just to deliver a single business-critical application to users and customers across a wide variety of environments. When there really wasn't an option because solutions weren't unified, weren't contextually aware, and were basically just a bunch of point solutions cobbled together with yet another management system. When you had to choose between sprawl or nothing.

It was a tough choice because there really was no choice at all.

Remember when you had to re-architect your network in order to add security solutions? And in the end you went the route of a configuration that would merely let you see attacks occurring, but weren't necessarily configured to stop them? All so you could maintain performance of applications and not risk violating service-level agreements.

Remember when you had to re-architect your network to include yet another device that operated only on WAN links? A solution that was a jack of all trades in the acceleration of applications, but a master of none because it was so focused on the link and network layers that it just didn't really dig into the application and improve its performance? Remember when the choice was either symmetric or asymmetric acceleration, because your budget wouldn't allow you to choose to do both? When choosing a symmetric solution meant you forfeited reliability or integrated security, but choosing asymmetric meant losing out on better performance over the WAN and other high-latency links?

Remember when your choice limited you to optimizing only two or three layers of the network stack with any given solution? When it was nearly a requirement to deploy one solution per layer of the stack just to ensure the applications were secure *and* fast? When you were forced to deploy solutions in such a way that you lost the agility offered by one solution because of the implementation of another? Because they weren't unified, they weren't collaborative, they weren't even remotely capable of communicating in a way that was optimized and maintained flexibility of your architecture.

Those were the days of niche application delivery products; of solutions with complex configurations aimed at addressing specific pain points like insecure platforms and applications, poorly performing WAN links, and less than optimally performing internal networks. When the need for multiple application delivery solutions were always a speed-bump in your well-performing infrastructure as well as your application deployment processes. When deploying application delivery solutions meant that you were introducing choke points into the infrastructure and reducing the capacity of existing application delivery network solutions, ultimately impeding performance, increasing the costs of management, and hard-wiring flexibility right out of your architecture.

Those were painful days, when the choices you had were really no choice at all.

Aren't you glad those days are (almost) over?