nemertes RESEARCH
Independence. Integrity. Insight.

2017

Securing the Internet of Things: Saving with Surge

Simplify Operations While Driving IT Savings

Waves of new networked devices are flooding enterprise networks, each a potential point of attack and breach. Most organizations use familiar security technologies to secure them: VLANs, ACLs, firewalls, Network Access Control. Nemertes studied deployments of Avaya Surge for securing such devices. It combines the virtues of a microfirewall with a centralized, policy driven holistic management system to reduce IT hands-on time in deploying and moving equipment and reduce spend on firewalls, NAC, and elsewhere.

Compass Direction Points:

- **Medical devices are multiplying—and moving.** Participants had thousands to hundreds of thousands of devices they ultimately need to secure; devices that move 302 times a year on average. They averaged 2,533 devices in their initial Surge rollouts.
- **Moving devices is expensive.** On average, $174 per move.
- **Surge can reduce per-move IT costs to $0.** IT no longer has to be present for or involved in moves—devices can simply be plugged in to the network in their new location, by anyone, and have all necessary security follow them.

By John E. Burke
CIO and Principal Research Analyst
Nemertes Research

Table of Contents

Table of Figures

Executive Summary

Organizations are experiencing waves of new kinds of devices flooding onto their networks, and use a variety of means to secure these new networked devices on the LAN. Using VLANs, router ACLs, firewalls, and network access control systems, IT departments try to strike the best balance among cost, complexity, and risk management. In the end, most methods end up requiring too much staff time to implement. When networked devices were few, and each device was typically placed and then left in place, the staff time consumed in placing it and moving it was relatively unimportant. As devices multiply, and flood onto the network, and as ever more of them are made mobile, the per-device time burden swiftly becomes insupportable.

Avaya Surge, by taking a radically different approach to securing the equipment's network access, can up-end cost/benefit calculations by eliminating much of the hands-on labor of IT, simplifying design problems, and sidestepping capacity management pitfalls. Per-device IT staff time to deploy a piece of equipment and later to move a piece of equipment drop to zero; planning times are greatly reduced as well. By so dramatically reducing operating costs while also simplifying management and improving security, a solution like Surge can make the new age of networked and mobile devices survivable for any IT department.

The Issue

Healthcare providers are at the leading edge of organizations facing a dilemma many other kinds of organizations (manufacturers, hospitality companies, universities) will soon face also: how to secure expensive, often movable, equipment on the network.

Securing Equipment the Old Way

Organizations typically take one of four approaches to securing their equipment:

- Put it on a secured VLAN or set of VLANS
- Put firewalls between secure VLANs and the rest of the network
- Put it on secure VLANs with Network Access Control (NAC) applied
- Put microfirewalls on each piece of equipment

Each approach has its merits (usually, familiarity) and drawbacks (usually cost and the amount of hands-on IT labor involved).

VLANs Alone

This is in some ways the simplest approach: define secure VLANs for protected equipment, and use router ACLs to limit access to them. No additional software is required, no additional hardware, just switches and routers and their in-built VLAN and ACL mechanisms.

It doesn't usually stay simple, though. In practice, this approach generally requires creation of multiple equipment VLANs, to segment different categories of equipment with different protection requirements from each other. This inevitably creates complication: to which VLANs can a given network port provide access? The more broadly useful a port is, the more care has to be exercised in mapping a piece of equipment's MAC address to the correct VLAN, both to ensure that what needs to talk to it can, and to ensure that what shouldn't see it can't. The more narrowly useful a port is, the more ports there have to be in each space. Many hospitals and clinics utilize a color-coding scheme, with each room or public space having some number of "red ports" and "green ports" and so on, and pieces of equipment tagged as to which color LAN they go on. Running out of ports of one color means no more pieces of equipment associated with that LAN can go in that space. Damaging a port in a given color might render the space unusable until the port is repaired; for example, an operating room might need four ports capable of handling monitoring equipment to be considered usable, so damaging a port might take it out of service until the port was repaired. To minimize such outages, the facility can pull extra ports in each room, knowing full well that they are putting in place more ports than they will use, most of the time, with added cost both in wiring and in switching, and having to balance the cost of added ports against the probability of failures—and against their ability to predict future growth in needs on each class of port.

For some classes of device, the secured VLANs would ideally propagate all the way back to the data center systems that interact with them. The more VLANs there are, the more complex and burdensome—and therefore the less likely—this becomes.

VLANs Plus Firewalls

These environments inherit all the same issues as the "VLANs alone" solution, but add to them the cost and management overhead of the more robust segmentation provided by firewalls. Different security organizations will define their needs for segmentation differently, but at a minimum will have to deploy firewall failover pair (for business continuity) per site. Some partition their segmentation further, deploying a pair per VLAN, for instance, or a pair per major section of a facility (e.g. a pair for emergency services, a pair for operating theaters, etc.).

The tradeoffs in the firewall scenario balance reducing the number of firewall pairs against cost, complexity, and concentration of risk.

Cost

Adding firewalls, of course, increases the cost of the solution; both capital expense for the boxes and operating costs for ongoing licensing and maintenance. Beyond that, though, there is increased administrative burden on IT to manage the additional pairs of devices, and each pair is typically administered separately rather than as a holistic security service. More hours spent managing and maintaining firewalls means fewer hours for everything else—and security staff, of whom there are never enough, always have a lot of "everything else" hanging over their heads and falling off the edges of their plates.

Complexity

Complexity in managing VLAN security is always increased by adding VLANs, classes of equipment, and firewalls. Which architecture the organization chooses shifts where the complexity is focused. Having a single firewall pair for all VLANs and zones means fewer to manage, but makes the rules on that one more complex. Having more firewalls (per VLAN or per zone) makes for simpler rulesets per firewall but more complexity in their deployment and management.

Concentration of Risk

Having a single pair of firewalls puts all services at the mercy of that pair, and anything that takes them out of service renders all the equipment useless. This puts an even higher premium on continuity of service than usual, and so leads to over-provisioning—buying more expensive gear with higher capacity and better built-in redundancies than might otherwise be needed.

To spread the risk out more broadly, IT can put many pairs of firewalls out and partition their work. Balancing how best to partition the work so that risk is minimized but costs do not become unbearable is a non-trivial exercise, and no single answer will serve everyone. Costs for the security infrastructure, both capital and operating, will rise, inevitably; but reducing the risk of broad service outages has several benefits:

- Improved ability to provide care and positive patient outcomes
- Reduced risk of death, injury, or worsened health
- Reduced risk of lawsuits and damages
- Reduced risk of regulatory infractions, fines and penalties
- Increased billable service hours for things like operating rooms

NAC

NAC systems take a more robust approach to securing movable devices by (a) knowing and recognizing each device and (b) dynamically assigning it to the correct VLAN when it is put on the network. This eliminates the complexity of switch-level mapping of gear to VLANs, one by one and manually, and keeps IT out of the loop when a piece of equipment moves. However, the VLANs and ACLs still have to exist—NAC is network *access* control only, and it doesn't build the necessary network segment to admit things to nor get traffic securely to the other end—so all the management overhead of the VLANs and ACLs will also still exist.

The downsides of NAC include the additional capital and operating costs, of course, as well as separate administration. Costs can be quite high, and are based on the number of entities tracked. Large healthcare environments may have a quarter of a million pieces of equipment transitioning from un-networked to networked, and so the expense can be considerable, and subject to steady growth over time.

Moreover, NAC has a significant limit on applicability in the context of non-traditional devices. Not all IoT-style networked devices support the functionality required to use NAC systems (typically the ability to act as an 802.1x supplicant), so a NAC system cannot be a complete solution.

Microfirewalls: An Emerging Option

A recent innovation in the space is to put a firewall in front of each piece of equipment, a so-called "microfirewall" built specifically for small-scope deployments, and to consume little space or power.

Microfirewalls generally remove the need for secured VLANs, since the secure segment is basically just the cable attaching the equipment to the firewall. They also minimize the complexity of the rulesets, since each firewall only has to know how to protect the one piece of equipment behind it. Where the VLAN was supposed to stretch all the way to the system back-end in the data center, though, the secure VLAN might still be needed, in which case the microfirewall approach has not succeeded in significantly simplifying network management.

Microfirewalls increase management complexity and overhead dramatically, because each one is managed separately. If there is a change to the protection profile for a given class of equipment, the firewalls for every instance of that type of equipment have to be touched, to update the policy. Patching and upgrading code has to be managed for each, as well. Since each is configured separately, there is a huge potential for misconfiguration, given how many there are, and increased misconfiguration means added risk of breach.

Surge Ahead

Avaya's Surge solution is aimed at ameliorating or eliminating many of the pitfalls of the various methods currently in use, while allowing a fast onboarding process for new equipment classes (via easy profile development) and for new pieces of equipment as well as IT-free device portability.

Like microfirewalls, Surge eliminates the need for separate VLANs by putting protection—the Open Network Adapter (ONA)—in front of each piece of equipment. Like NAC, Surge recognizes the equipment the ONA is associated with and automatically applies the right protection profile to it, wherever it is plugged in. Like firewalls and microfirewalls, ONAs are security devices first and foremost, so by default they implement an approach of "deny-all, allow only by policy."

However, unlike microfirewalls, Surge uses a centralized management console to provide policy-driven, holistic management of the entire Surge deployment—no device-by-device configuration.

And, unlike any of the other solutions, Surge is capable of establishing a virtual, overlay network—an encrypted tunnel, in essence—directly from an ONA to systems in the data center, enabling end-to-end security partitioning where needed.

Studying Costs Before and After Surge

Nemertes conducted interviews with four early users of Surge to identify and quantify the benefits they have realized by deploying it.

The Participants

Participants were large healthcare providers, averaging:

- 69,000 employees
- 905 locations
- $13.55B annual revenues

They were in different geographies from each other, and each participant's environment comprised a mix of hospitals, clinics, and medical office facilities. All conversations took place in January of 2017.

The Savings

Although all are in the early stages of deployment, the three that have gotten past the proof of concept phase have seen solid benefits and are proceeding with broader rollouts. In terms of IT staff time spent on putting equipment on the network:

- All report hands-on IT time per placement of a new device dropped to 0
- All report hands-on IT time per relocation of a device dropped to 0
- All expect to improve utilization of the protected resources, improving patient outcomes and driving higher revenue
- All anticipate significant hardware or software savings over their current solution

Cost Factor	Average Reduction	Maximum Reduction	Improvement Factor (Max)
IT time, initially place protected device	100%	100%	
IT time, protected device move	100%	100%	
IT time, profile development	48%	90%	10x
IT time, profile review	34%	75%	4x
Wiring/infrastructure costs	20%	60%	2.5x
Firewall costs	24%	97%	30x

Figure 1: Representative time and cost savings with Surge

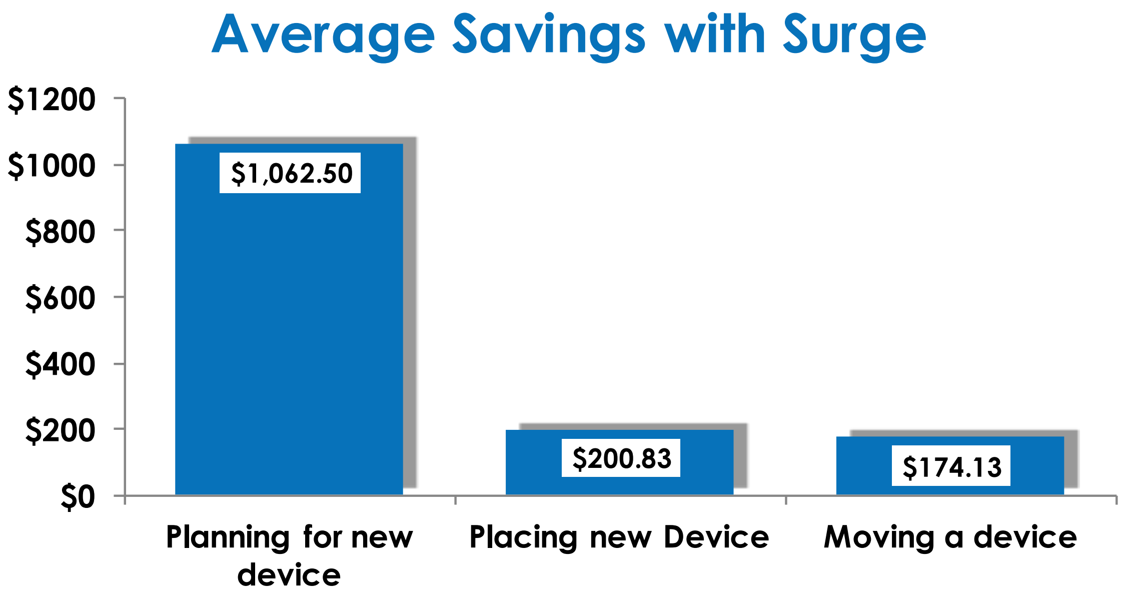

Figure 2: Average Surge-Driven Savings

	Per-Instance Savings (Avg)	Instances per Year (Avg)	Expected Average Annual Savings
Deploying a device	$200.83	2766	$555,430
Moving a device	$174.13	833,836	$145,198,631
Total			$145,754,061

Figure 3: Annualized Average Estimated Surge-Driven Savings for Participants

Participant 1: Saving with Surge vs. VLANs

This participant, using 802.1x authenticated network access and VLANs, expended about an hour placing a new piece of equipment on the network, and only about six minutes per move per piece of equipment. However, given the volume of moves they were facing in the environment they decided to move to Surge, moves there, alone, consumed an entire FTE's worth of time pre-Surge.

Post-Surge, prepping a device to deploy took IT less than a minute, and again, non-IT staff were able to perform the actual placement and connection of equipment, as well as all equipment moves.

Pre-Surge, IT would spend 24 hours of senor staff time planning the security profile for a new type of device. Post-Surge, they are spending eight hours, a savings of 66%. However, this organization's biggest savings is coming with the ability to break away from their system of multiple color-coded sets of VLANs and LAN ports in operating rooms and elsewhere. The ability to use generic LAN ports without compromising security is going to reduce their related infrastructure costs by 60% by eliminating redundant cable pulls, patch panels and cables, and unnecessary switch ports.

Beyond that, this organization expects to garner top-line benefits as well: by reducing the incidence of facilities being put out of service due to problems with availability of the right mix of secure VLAN ports—a loss of revenue as well as an IT cost to fix the problem—the hospitals and clinics will be able to use more rooms, more of the time, for patient care.

Participant 2: Saving with Surge vs. VLANs+Firewalls

This participant expended eight hours of IT network staff time to deploy a new piece of equipment prior to Surge, and another eight on most equipment moves. Prior to Surge deployment, in the converted portion of the environment this amounted to two FTEs worth of time per year devoted solely to connecting equipment to the network.

With Surge, they expend no significant IT staff time on typical deployments or moves now—a 100% reduction in hands-on IT time installing the medical equipment. Facilities or clinical engineering staff is able to plug equipment in, and

Surge does the rest. Behind the scenes, IT spent less than a minute per device associating each with its ONA and assigning it to a security profile.

Pre-Surge, IT spent 20 hours of senior staff time planning the security profile for a new type of device. Post-Surge, they are spending 2.5 hours, a savings of 87.5%.

More importantly, though, with Surge in place they will be able to radically downsize their firewall network. In order to reduce the impact of outages and configuration problems, as well as to partition the work of the firewalls and reduce the complexity of rulesets, they have adopted a broadly-distributed firewall model: a firewall pair per physical section of site, with as few as five and as many as 60 sections per site. That means anywhere from five to 60 failover-pairs of firewalls per site. With Surge, they are able to reduce to one pair per site without re-concentrating risk or adding complexity, for a reduction of 80% to 96.67%.

Participant 3: Saving with Surge vs. NAC

This participant will see per-device deployment times drop from 1.5 hours of network staff time per device down to zero, as the others are seeing, and no-IT-involvement movement of devices. And, although they are expecting to realize savings on their NAC solution by gradually removing devices from it (and thereby lowering the annual licensing) they are not eliminating it since they use it for regular computing gear.

For this organization, another aspect of savings is deferring replacement of networked medical devices that, though still able to provide their medical function, are no longer being supported on the network side (with security patching, for example). They anticipate using Surge to secure these devices well enough that they can continue to run them until they are no longer fit for their core medical purpose.

Conclusions and Recommendations

Clearly, there are many ways to secure equipment on a LAN, and just as clearly, most of them entail a lot of effort and expense to maintain. Avaya Surge, by taking a radically different approach to securing the equipment's network access, can up-end the cost/benefit calculations by reducing or eliminating much of the hands-on labor of IT, simplifying design problems, and sidestepping capacity management pitfalls. In an environment where more kinds of equipment and more pieces of equipment want to be on the network every year, any solution that can so dramatically reduce costs and simplify security operations while improving security and opening up the possibility of real top-line benefits to the organization deserves a close evaluation.

About Nemertes Research: Nemertes Research is a research-advisory and consulting firm that specializes in analyzing and quantifying the business value of emerging technologies. You can learn more about Nemertes Research at our Website, www.nemertes.com, or contact us directly at research@nemertes.com.