Network Security in the Age of *Any* Computing

Risks and Options to Control Mobile, Wireless, and Endpoint Devices

Version 1.2

Released: April 1, 2011

Securosis, L.L.C. 515 E. Carefree Blvd. Suite #766 Phoenix, AZ 85085 **T** 602-412-3051 info@securosis.com www.securosis.com

Author's Note

The content in this report was developed independently of any sponsors. It is based on material originally posted on the Securosis blog <http://securosis.com>, but has been enhanced, reviewed, and professionally edited.

Special thanks to Chris Pepper for editing and content support.

Sponsored by ForeScout Technologies

ForeScout Technologies is a leading provider of automated security control solutions for Fortune 1000 enterprises and government organizations. With ForeScout, organizations can accelerate productivity and connectivity by enabling people to access corporate network resources where, how and when needed without compromising security. ForeScout's CounterACT platform for network access control, mobile security, threat prevention and endpoint compliance empower access agility while preempting risks and eliminating remediation costs. Because ForeScout's solutions are easy to deploy, unobtrusive, intelligent and scalable, they have been chosen by over 1000 of the world's most secure enterprises and military installations for deployments spanning 37 countries. For more information, visit www.forescout.com

Copyright

Table of Contents

The Risks of *Any* Computing

Everyone loves their iDevices and Androids. The computing power that millions now carry in their pockets would have required a raised floor and a large room full of big iron just 25 years ago. But that's not the only impact we see from this wave of *consumerization,* the influx of consumer devices requiring access to corporate networks. Whatever control you thought you had over the devices in the IT environment is gone. End users pick their devices and demand access to critical information within the enterprise. Whether you like it or not.

And that's not all. We also have demands for unfettered access from anywhere in the world at any time of day. And though smart phones are the most visible devices, there are more. We have the ongoing tablet computing invasion (iPad for the win!); and a new generation of workers who demand the ability to choose their computers, mobile devices, and applications. *Even better, you aren't in a position to dictate much of anything moving forward.* It's a great time to be a security professional, right?

Whatever control you thought you had over the IT environment is gone. End users pick their devices and demand access to critical information within the enterprise. Whether you like it or not.

Sure, we could hearken back to the good old days. You know — the days of the BlackBerry, when we had some semblance of control. All mobile access happened through your BlackBerry Enterprise Server (BES). You could wipe the devices remotely and manage policy and access. Even better, you owned the devices so you could dictate what happened on them.

Those days are over. Deal with it.

Risks

We call this concept ***any*** computing. You are required to provide users with access to critical and sensitive information on *any* device, from *any*where, at *any* time. It's scary as hell.

To be clear, any device really means *any* device. But we will focus on mobile devices in this paper. That's for obvious reasons — particularly because we have had a long time to figure out how to deal with traditional computer endpoints. These consumer mobile devices are a new ballgame. So we need to get a handle on

them, and we need it now. Let's start by taking a step back and quickly examining the risks. If you want more detail, check out our white paper on Mobile Device Security (PDF):

1. **Lost Devices:** Way too many numb-nuts you work with manage to lose laptops, so imagine what they'll do with these much smaller and more portable devices. They will lose them, and your sensitive corporate data on them. But that's not the only risk. You need to be wary of device sales — folks often use their own the devices, access your sensitive data, and eventually sell the devices. A few of these people will think to wipe their devices first, but you cannot rely on their memory or sense of responsibility.

2. **Wireless Shenanigans:** Pretty much every computing device for sale includes a WiFi radio, which means folks can connect to any malware-laden public network. And they do. So you need to worry about what they are connecting to, who is listening (man in the middle), and whatever else can interface with network connectivity. And rogue access points aren't only in airport clubs and coffee shops. The odds NetStumbler could find some 'unauthorized' networks in your own shop are pretty good. Even worse, plenty of folks use 3G cards to get a direct pipe to the Internet — bypassing your egress controls completely, and if they're generous they might provide an unrestricted hotspot for their neighbors. Did I hear you say ubiquitous connectivity is a good thing?

3. **Malware:** Really? To be fair, malware isn't much of an issue on phones now. But you can't assume it never will be. More importantly, consumer laptops may not be protected against today's attacks and malware. Even better, many folks jailbreak their devices to load the new shiny applications — not noticing that in the process they disable many of the built-in security features. Awesome.

4. **Configuration:** Though not necessarily a security issue, you need to consider that many of these devices are not configured correctly. They load applications they don't need and turn off key security controls, then connect to your customer database. So *any* computing creates clear and significant management issues as well. If not handled correctly, these exponentially expand the attack surface.

Lots of folks want to tell you all about how their new widget or yet another agent on all your devices will protect you from all of the ills of society. And it is important to look at device level security; but in this paper we look at these issues from a network-centric perspective. Because it's all about control. You don't control the devices, especially devices belonging to guests and contractors visiting your IT infrastructure, so you need to work at a layer you do control — **the network**. So we will examine a few network architectures to deal with these devices. We will also look at some network security technologies — most notably Network Access Control (NAC) — that can help protect critical information assets from the onslaught of mobile devices.

Business Justification

Finally, let's deal with the third wheel of any security initiative: business justification. Ultimately, in order to deploy any new security technology, you'll need to make the case to management that it's a must-have. We all know how often you get "nice to have" projects over the finish line. Of course, you could default to the age-old justification of fear — wear them down with all the bad things that *could* happen. But with *any* computing it doesn't need to be that complicated because the threats are real.

1. **List top line impact:** First we need to pay attention to the top line, because that's what the bean counters and senior execs focus on. So map out what new *business* processes these devices enable, and get agreement that their top-line impact is bigger than a breadbox. It will be hard or impossible to forecast true revenue impact, so the goal is to get acknowledgement that positive business impact is real. In the field, we see mobile devices rolled out for a good reason: they result in more business. So leverage that.

2. **New attack vectors:** Next have a very unemotional discussion about all the new ways to compromise your critical information via these new processes, running on these mobile devices. For example, you can enumerate a few attack vectors made possible by guests, contractors, rogue mobile devices, and rogue access points run amok. Again, you don't need to throw FUD (fear, uncertainty, and doubt) bombs, because you have reality on your side. *Any* computing does make it harder to protect information.

3. **Close (or not):** Basically you are in a position to now close the loop and get funding — not by selling Armageddon, but instead providing a simple trade-off. The organization needs to support *any* computing for lots of business reasons. That introduces new attack vectors, putting critical data at risk. It will cost $X to more adequately protect the information. Yes, you could talk about downtime and customer service and brand. But ultimately, we believe that business justification is about *objectively* presenting both sides of the story and allowing the business folks to make a business decision.

Containing Access

Now that we understand the risks of these mobile devices, we need to start thinking about how to protect the network and therefore the data they are accessing. Of course, we have controls at the device layer to protect them and ensure proper configuration. But in this paper we will think about how to architect and secure the *network* to protect critical data. The first step is restricting access to key portions of your network to only folks who need it.

Segmentation is your friend

There is an old saying, "out of sight, out of mind," which could be rephrased for information security as, "out of reach, out of BitTorrent." By using a well thought-out network segmentation strategy, you can keep the critical data out of the clutches of attackers. Okay, that's an overstatement, but segmentation is the first line of defense to protect key data.

You want to make it as hard as possible for the data to be compromised, and that's why you need to put up as many obstacles as possible in the path of attackers.

You want to make it as hard as possible for the data to be compromised, and that's why you need to put up as many obstacles as possible in the path of attackers. Unless you are being specifically targeted, simply avoiding being the path of least resistance is a decent strategy. The fewer folks who have access to something, the less likely that access will be abused, and the more quickly and effectively we can figure out who is the bad actor in case of malfeasance. Not that we believe the PCI-DSS v2.0 standards represent even a reasonable low bar for security controls, but they do advocate and require segmentation of cardholder data.

Here is the specific language:

> *All systems must be protected from unauthorized access from untrusted networks, whether entering the system via the Internet as e-commerce, employee Internet access through desktop browsers, employee e-mail access, dedicated connections such as business-to-business connections, via wireless networks, or via other sources. Often, seemingly insignificant paths to and from untrusted networks can provide unprotected pathways into key systems. Firewalls are a key protection mechanism for any computer network.*

One architectural construct for segmentation is our idea of vaults, which really are just a different way of thinking about segmentation of *all* data — not just cardholder data. This entails classifying data sources into a few tiers of sensitivity and then designing a control set to ensure access to only authorized parties.

The goal behind classifying critical data sources is to ensure access to network-based resources is only provided to the *right* person, on the *right* device, from the *right* place, at the *right* time. Of course, that first involves defining rules for who can come in, from where, when, and on what device. And that effort is time consuming and difficult. But defining these granular policies needs to be done.

Once the data is classified, user access scenarios are evaluated, and the network is segmented accordingly, we need to authenticate the user. An emerging and flexible means of enforcing access to only those authorized devices is to look at something like risk-based or adaptive authentication, where the authentication isn't just about two or more factors, but is instead dynamically evaluated based on any number of data points — including *who* you are, *what* you are doing on which device (factoring in the security posture of the device), *where* you are connecting from, and *when* you are trying to gain access. And given the legitimate questions regarding token-based authentication in the wake of the RSA breach, we need to make sure authentication is based on more than just what you have (the token).

The obvious weakness of any control structure focused purely on initial authentication is that a device can be compromised after entry — and then all the network controls are irrelevant because the device already has unfettered access.

This certainly works well for ensuring only the right folks get in, *but what happens once they are in*? The obvious weakness of any control structure focused purely on initial authentication is that a device can be compromised after entry — and then all the network controls are irrelevant because the device already has unfettered access. A deeper look at risk-based authentication is beyond the scope of this research project, but warrants investigation as you design control structures.

Look very critically at how the network controls can be bypassed. If a machine is compromised after getting access, that is a problem unless you are constantly scrutinizing who has access on a continuous basis. Consider unauthorized physical access to your network. That could be a hijacked physical port or a rogue wireless access point. Either way, someone gets physical access to your network and bypasses the perimeter controls.

Architectural Straw Man

Now let's talk about one architectural construct in three different use models for your network, and how to architect a network in three segments, depending on the usage model for access.

1. **Corporate network:** This is for staff (or fully trusted contractors) who have physical access. Either via wired connections or wireless access points.
2. **External mobile devices:** These devices access corporate resources via *uncontrolled* networks. That includes home networks, public wireless networks, cellular (3G) networks, and even partner networks. If your network security team can't see the entirety of the ingress path, then you need to consider it an external connection and device. Better safe than... well, you know.
3. **Internal guest access:** These are devices that just need to access the Internet from inside one of your facilities. Typically these are smartphones, notebooks, or tablets used by employees or guests, but we also consider businesses (retail/restaurants, healthcare facilities, etc.) providing access as a service to their customers.

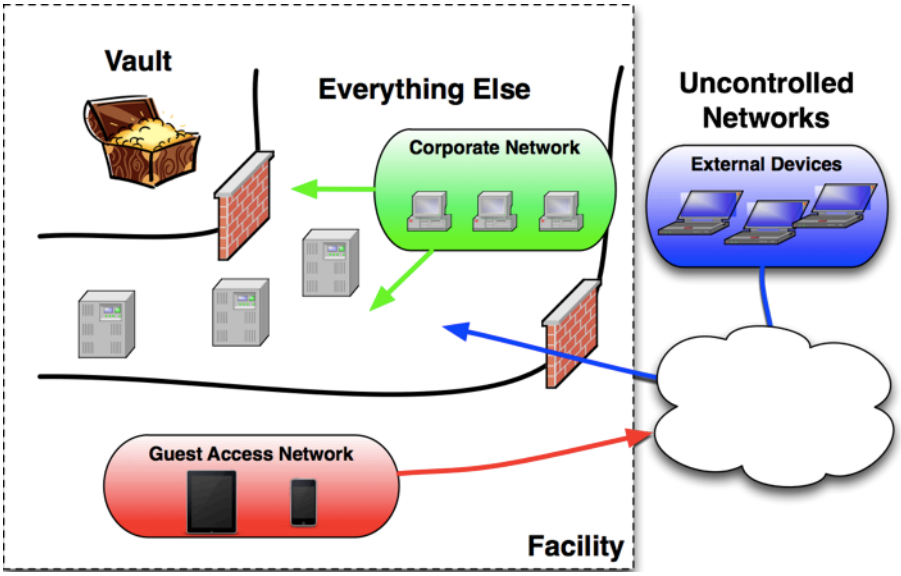

We want to provide different (and increasing) numbers of hoops for users to jump through to get access to important network-based resources and the associated data. The easiest to discuss is the third case (internal guest access), because you only need to provide an egress pipe for those folks. We recommend *total physical isolation* for these devices. That means a totally separate (overlay) wireless network, with a different uplink to the Internet. Yes, that's more expensive. But you don't want a savvy attacker figuring out a way to jump from the egress network to the internal network. If the networks are totally physically separate, you eliminate that risk.

The techniques to support your corporate network and external mobile devices are largely the same under the philosophy of “trust, but verify.” So we need to design control sets to scrutinize users. The real question is how many more hoops do you put in place for an external device than one with physical access to your network. And how do you classify the data to establish tiers of sensitivity. This feeds directly into how you define your access policies.

Defining Policies

Given the dynamic nature of business, access policies should provide sufficient flexibility to meet business needs. We are big fans of models to understand your exposure points and choose what needs to be protected and when, so first map out the access models you need to support. With that information you can design policies to provide users with the right access at the right time from the right device.

To illustrate, let's look at a few categories of mobile devices, since they are the poster children for *any* computing and typically the hardest to secure. First we will define three general categories of mobile devices trying to connect to your network:

1. **Corporate devices:** You have issued these devices to your employees and they are expected to get full access to pretty much whatever they need. You'll want to verify both the user (with strong authentication) and the device itself, including configurations, security settings, etc. It is also important to monitor what the device is doing to ensure authorized use after authentication and connection.

2. **Personal devices:** Sure, it's easy to just implement a blanket policy of no personal devices. There are big companies doing that right now, regardless of user grumpiness over not being able to use their fancy new iPads at work. But if draconian isn't an option in your shop, you could move authenticated, unauthorized devices onto a logical network configured only for outbound Internet access. Alternatively, an option may be providing access to non-critical resources (employee wikis and the like), while blocking access to corporate email servers, assuming you don't want company email on these devices.

3. **Everything else:** Plenty of guests show up at your facilities and try to connect to your networks — both wired and wireless. If they successfully gain access via WPA2 or a physical port, you need to decide whether to bounce them from the network or limit their access to network resources, depending on where they can get to and what kind of guest they are. Another option is to have them register to be able to access the network. This represents the 'access' part of Network Access Control.

Depending on your pain threshold, many other device types and usage models can be profiled to create specific enforcement policies. Granularity is only limited by your ability to map use cases, design access policies for your requirements, and technically verify and enforce these policies on the appropriate devices. Don't forget you can also implement policies based on roles. For instance, your marketing group might have network access with iPads, since every good marketer needs one. But if engineers do not have a business

justification for iPad use, that group could be blocked. Policies aren't defined merely by what device the user has, but also on who they are and how their device is configured.

Posture-based Policies

What about policies based on defenses implemented on the endpoint or mobile device — such as AV, full disk encryption, and remote wipe? Clearly you define additional policies based on the defenses in place on the devices as well. For instance, restricting users without certain patches on their device is legitimate. Or you might want to keep end users off your protected network segment if they don't have full disk encryption active, to avoid breach disclosure if they lose the device.

Policies aren't defined merely by what device the user has, but also on who they are, and how their device is configured.

It's not just about knowing what the device is, and who is using it, but also what's on it. As you can see, this problem includes at least 3 dimensions, which is why getting policies right is a prerequisite for controlling access. We'll talk more about improving policies incrementally later.

Which, once again, brings up our main point. Make sure you can enforce security policies that reflect your desired security posture, given the context of your business processes. Don't force your security policy to map to your enforcement mechanisms.

Enforcement

Now that we have defined policies, it's time to enforce them on all devices accessing your stuff, especially those pesky mobile devices. We'll break this section up into groups of enforcement technologies such as NAC (Network Access Control), firewalls, and other network layer controls (such as VLANs and physical segmentation). Each technology has pros and cons. There are no 'right' answers — just sets of compromises that must be weighed against the various alternatives. And to be clear there are no 'either/or' decisions here. These technologies work hand in hand — for example, VLANs are typically used to implement policies defined on NAC devices.

NAC

Yes, we are talking about that NAC — you know, Network Access Control. No, it's not a dead technology. In fact, as you'll see, it's a key method for restricting access by mobile devices and other endpoint devices. Let's take a look at what NAC really does. Basically NAC scrutinizes the devices on your network to make sure they are configured correctly, and accessing what they are supposed to, as defined by additional policy attributes such as user, location, etc. That's the *access control* aspect. The initial use case for NAC was guest/contractor access, where it was particularly important to ensure any devices connecting to the network were authorized and configured correctly.

Given the requirement to ensure the right devices access only the right resources, integrating mobile device security with Network Access Control offers a means to restrict access to the critical information stores, so you know these devices are legitimate.

Of course, under this *any* computing concept — driven by mobile devices — **every** device now should be considered a *guest,* whenever feasible. Given the requirement to ensure the right devices and users access only the right resources, integrating mobile device security with Network Access Control offers a means to restrict access to the critical information stores, so you know these devices are legitimate. Which is what it's all about.

So what's the issue? Like most interesting technologies, there are issues of both functionality and perception with the technology. Some offerings are very complex requiring multiple devices and significant configuration and integration, hindering mass-market deployment. Though depending on the depth of the problem and the required degree of integration with existing infrastructure, this may

not be an issue. After 10 years, many NAC products have worked through the kinks of enterprise deployment.

There are also some sticky deployment design issues, depending on how you want to restrict access. Firstly NAC devices need to see most of the traffic flowing through your network, which requires them to be highly scalable, and if deployed in-line requires sensitivity to latency in the deployment design. You must be able to scan traffic and detect issues in real time or very close to it.

As folks who follow markets for a living we know that once the hype around any market starts dying down, the technology starts becoming more prevalent — especially at the large enterprise level. NAC is no different. You don't hear a lot about it, but it's happening, largely driven by the need to manage risks associated with the explosion of distinct user types, variety of network access options, the proliferation of mobile devices and the need to more effectively segment networks.

Firewalls

As described above, the PCI Guidance specifies firewalls as key for network segmentation and protection of sensitive information. We agree that front-ending sensitive data stores with firewalls is a best practice to control access to sensitive segments. Unfortunately, traditional firewalls only understand IP addresses, ports, and protocols. With many newer web-based applications — increasingly encapsulated on port 80 — that can be problematic.

This is driving the evolution of the firewall to become much more application aware. We have covered that evolution extensively, so we won't rehash it all here. But when using network segmentation for securing mobile devices, scalability of application layer inspection remains a major concern. These devices need to inspect and enforce application layer policies at multi-gigabit internal network speeds. That's a tall order for today's firewall devices, but as with everything else in technology, devices continue to mature and get faster. All hail Moore's Law!

These evolved firewalls are also instrumental for implementing network segmentation architecture to support *any* computing.

Network Layer Controls

The path of least resistance tends to require leveraging devices already in place, which means using the built-in capabilities of network switches and routers to enforce required segmentation and access control capabilities. First let's examine the brute force approach, which is physical segmentation. As we described, we believe Internet access for mobile devices and guests should be on a totally disparate network. You don't want to give a savvy attacker any chance to jump from you guest network to your internal net. This level of physical segmentation is great when the usage model supports it, but for most computing functions it doesn't.

So many folks leverage technologies such as VLAN (virtual LANs) to build logical networks on top of a single common physical infrastructure. At the theoretical level this works fine, and it will likely be enough to pass your PCI assessment. That said, the objective isn't to get the rubber stamp, but to protect the information. So let's take a critical look at where VLANs can be broken and see whether that risk is acceptable.

There are many ways to defeat VLAN-based segmentation, including VLAN hopping. To be fair, most modern switches can detect and block most of these attacks *if configured correctly*. That's always the problem — devices are only as strong as their configuration, and it's never safe to assume a solid and secure configuration. Nor is leaving the security of your critical data to a single control. Layers of security are good; more are better. But everyone already has switches, which support VLANs and physical segmentation, so this will continue to be a common mechanism for restricting access to sensitive data, which is a good thing. Yet keep in mind the challenges of administering policy across various network and security infrastructure.

Which brings us to the main point of this enforcement concept — it's not an either/or situation. The answer is all of the above. VLANs + firewalls + NAC provide a comprehensive, layered system to ensure only the right devices and users access critical data. That doesn't mean you need to do everything, depending on the real sensitivity of the data, but it doesn't hurt.

The answer is all of the above. VLANs + firewalls + NAC can work together to provide a comprehensive system to ensure only the right devices are accessing critical data.

Device Health

As described in the Mobile Device Security paper (PDF), there are many risks — especially in light of the prevalence of malware on traditional PCs and misconfiguration on mobile devices. When evaluating network layer controls, assess your need for device context — the security posture of any device connecting to your networks. Obviously the bar will be set differently depending on what the device is accessing.

So, for example, you may need to carefully scrutinize any device accessing your 'vault' of protected cardholder data. In that case, NAC + firewall is critical. But for everything else, a VLAN-only structure may be sufficient — even though you'd rather not have any general-purpose applications compromised — if it's a manageable risk. And depending on the deployment architecture of your NAC gear, you may be able to inspect all your enterprise traffic and restrict access accordingly.

The point is to design network security controls based on what they protect, and it's okay to have different tiers of protection based on what resources are available on that particular segment (whether physical or logical).

Visibility vs. Control

One other consideration for enforcement is whether you want to focus on knowing what's happening in your environment (visibility), or need to actively block or limit traffic that violates policy (control). Actually, you need to do both, but the question is when. This has a profound impact on the architecture you implement initially and how you evolve the devices over time. If your initial use case is to gain some visibility to see who is accessing what from where, you can send log files and/or traffic flows to another device for analysis. This more audit-centric use case is popular with folks who are unable to implement stronger controls (like blocking), due to deployment, latency, and/or political issues.

Another situation where visibility is helpful is to look for what we call "not normal" in your environment. This is out of character behavior, which may indicate of a compromise. By tracking the typical traffic flows from devices and then setting rules to alert on variations in activity (through SIEM or compatible NAC devices), you can pinpoint when you have an issue. Again, this kind of anomaly detection (pioneered by technologies like network behavioral analysis) can be very useful to help prioritize an investigation, as well as providing another data point to build policies on. So if a mobile device that typically only connects to the email server suddenly starts probing key financial file stores, you can take a more active remediation approach.

Active Remediation

If you want ultimate control over the traffic that reaches these sensitive segments; you need a box to inspect all traffic either in-line or via a span port; with the ability to block or limit offending traffic either directly (inline) or via configuring another device like a switch, router or firewall. We'll discuss integration with these devices next. There are multiple ways (including TCP resets, DNS tampering, and ARP poisoning) for passive devices to block traffic by intervening and tearing down offending sessions *without* being inline, but they all present security issues and should not be the only remediation mechanism in play.

There are no right or wrong answers — only considerations for designing your own network security controls.

Again, there are no right or wrong answers — only considerations for designing your own network security controls. Clearly you will be dealing with more and varied types of users and devices moving forward, and they need deeper access to more sensitive information. *We advocate a phased approach, where the first wave focuses on visibility.* This helps demonstrate the seriousness of the issue without risking interference with business. You know it's a bad day when you block something that results in the loss of a multi-million dollar deal.

The next step becomes implementing remediation controls to enforce network access policies, which means integrating with your active network/security devices.

Integration

Supporting *any* computing requires organizations to restrict access to specific communities of users/devices based on organizational policies, as discussed above. In order to accomplish this, you need to integrate with your existing installed base of security and networking technologies, ensuring management leverage and reducing complexity. No easy task, for sure. So let's discuss how you can enforce access policies while playing nicely in the larger sandbox.

Authentication

When an endpoint/mobile device connects to the network, you can start with either a specific authentication or network-based detection of the device, via passive monitoring of the network traffic or the MAC address of the connecting device. The choice of how strong authentication must be comes down to whether building policies based on device and/or location will be granular enough to enforce the access policies. If you want to take *who* is using the device into account in the policies, you will need to know the identity of the user. There are techniques to identify users passively, but we prefer stronger methods for determining identity, which require integration with an authoritative source for identity information. The integrated directory is typically be Active Directory, LDAP, or RADIUS. Authentication is via an agent (persistent or non-persistent), a connection portal (provided as part of the NAC solution), or a protocol such as 802.1X (which can create challenges on mobile devices, such as limited functionality supplicants).

Keep in mind that identity is a dynamic beast, and users & groups are constantly changing. So it's not sufficient to provide a one-time dump of the directory. You'll want to track user/group moves, adds, and changes; on an ongoing basis. Authentication time is also the right time to figure out what's going on with the device, which involves inspecting it to understand its security posture.

Endpoint/Mobile Device Integration

The first decision is how deeply to scrutinize endpoints/mobiles when they connect. Obviously there is a time factor to scanning and checking security posture, which can cause user grumpiness. Most organizations want to make sure devices are properly configured upon access, many aren't ready to react to the answers they may get. Do you block access when a device violates policy? Even when the user has a legitimate and business critical need to use the network? As we discussed above, you may want to define policies based on the security controls in place on the endpoints/mobiles.

Compromising your security by providing access to compromised devices makes no sense, so what remediation should happen? Do you patch the device? That requires the ability to integrate with a patch

management product. Do you reconfigure the device? Or update the endpoint protection platform? This depends on the nature of the policy violation and which information that user can access, but you want *options* for how to remediate. And each option requires support from your NAC vendor. Supporting mobile devices can be more complicated, since the technical ability to manage the mobile device varies – and that doesn't even factor in the policy/cultural issues of supporting guest users and personal devices. You *could* just ignore the details and block users with devices which don't comply with policy, but this tends to end with your rainmaker calling the CEO because she can't get into the ordering system to book that critical deal. Which rarely works out well for you. Limiting access may be a better option, but ultimately you need to balance usability versus security.

Another consideration is that devices can be compromised **after** connecting. Detecting compromised devices involves both re-authenticating devices periodically (to ensure a man in the middle hasn't happened), as well as assessing the security posture of the device on a periodic basis. Another tactic is to detect compromised devices by their behavior — which requires continuously checking devices for anomalous behavior. Most NAC devices are *already* monitoring the network to detect new devices, so this anomaly detection capability is frequently available, as we discussed above.

Another consideration is that devices can be compromised **after** connecting, including continuously checking devices for anomalous behavior.

Now that you know the posture of the endpoint/mobile, you can determine the appropriate level of access to grant it, enforcing that policy with an active network/security device via any of various integration techniques.

Network Integration

There are plenty of ways to enforce network access policies using switches and firewalls. Let's take a look at the major techniques:

1. **Inline device:** Obviously one option for enforcing access policies is to be in the middle of the connection and able to block unauthorized devices as needed. Networking infrastructure players who offer NAC can provide multipurpose boxes that act as hybrid enforcement points. At some times, such as during authentication, the devices are actively involved, but after authentication they just passively monitor activity unless a policy is violated. There isn't much more to say about it, but deploying a device inline has a major impact on network design.

2. **CLI:** The good old command line is still one of the more popular methods of integrating with active devices. This involves the NAC equipment establishing a secure, authenticated session (typically using

SSH or SSL) with a switch or firewall and making the appropriate changes. That might mean moving a user onto a guest VLAN or blocking their IP from accessing protected networks. Obviously this requires specific integration between vendors, but given that only a handful of vendors control the switch and firewall markets, this isn't too daunting. That said, there can be delays in compatibility when network or security gear is upgraded, so make sure to check NAC support before any upgrades.

3. **802.1X:** The 802.1X standard is used for authentication on connect (as described above), for which it is well suited. But the protocol also includes an option to send enforcement policies to endpoints, which gets far more complicated involving integration and support from a variety of additional devices. Even though 802.1X is a mature standard, interoperability can still be problematic in heterogeneous network/security environments, which is compounded when considering guest and employee-owned devices. Obviously using 802.1X for enforcement is a non-starter if only a limited function supplicant is available (typical for mobile devices). Individual vendors have generally sorted interoperability between their own NAC and general networking products, but making 802.1X work at enterprise scale is nontrivial.

4. **SNMP:** Another option for integration with switches is using SNMP to send updates to the networking gear. The advantages of SNMP clearly center around ubiquity of support, but security is a serious concern (especially with early versions of the protocol), so ensure you pay attention to device authentication and session security to ensure the security and integrity of the communications.

All of the above

As usual, there is plenty of religion about which integration technique is best, which continues to amuse us. Our stance hasn't changed: diversity in integration techniques is better than no diversity. We also prefer multiple enforcement tactics — multiple layered controls provide more hurdles for attackers making it harder to compromise the system. That means you want the ability to enforce access policies at the switch, perhaps with a different inline device (like a firewall) segmenting your most sensitive information.

You can also leverage some direct endpoint/device-based methods to mess with a user's IP stack, including measures such as TCP resets, DHCP tampering, and ARP poisoning to spike sessions or redirect users to an authentication portal. As discussed before, there are clear limitations to some of these protocol-based techniques for security. But depending on the use case for access and the sensitivity of the protected data, they may provide a good compliment to more active remediation options.

The Weakest Link

Finally, you need to balance deployment breadth against security, because the NAC device needs access to all traffic to enforce access control policies across the enterprise. If there is a blind spot on the network, that makes it easier for an attacker to gain access and compromise an authorized device — and then to bypass network-based access control. The weakest link is the link you don't see — do you see?

So starting with your proof of concept, focus on making sure you have sufficient visibility to meet security requirements. This often requires savvy use of SPAN ports for passive monitoring, and designing choke points for inline devices to avoid unacceptable latency for applications. We always recommend a focus on protecting the most critical stuff when planning a roll-out — be very sure to enforce Network Access Control where it's needed most.

Audit trail

We all saw *Terminator* and have no intention of allowing SkyNet to reconfigure our networks without approval, so focus on auditing all changes made to any network infrastructure equipment by NAC (or any other technology, for that matter). Not only is this standard operational practice to facilitate troubleshooting, it's also required for regulatory compliance. The assessor will need to be shown that you understand why any change is made to the network (especially to any protected segments) and have documentation of who made each change.

Monoculture

As Dan Geer so eloquently explained years ago, there are clear disadvantages to a single-vendor environment — to being dependent on a monoculture, which are aggravated if the vendor's goals are incompatible with your own. Of course, vendors who have made countless acquisitions over the past few years tell the other side of the tale, playing up the benefits of integrated management and enhanced integration in their unified environments. Unsurprisingly, we believe the truth lies in the middle.

Just because you buy products on the same purchase order doesn't mean they are integrated.

First you need to understand the true extent of integration within a vendor's product line. Just because you buy products on the same purchase order doesn't mean they are integrated. This is particularly acute for companies that grow via acquisition. If the integration is real, then evaluate the benefits of integrated management for your deployment use cases. If you have modest requirements, simple SNMP enforcement might be more than sufficient, and can work well in a multiple-vendor environment. But if you need 802.1X, interoperability may be problematic and overly complex. So there really is no simple answer to the question of single vs. multiple vendors.

To be clear, vendors that do not offer switches and firewalls cannot provide the entire infrastructure, so they aggressively support the infrastructure market share leaders. *The best tactic is to be skeptical — about everything.* Don't trust a slide deck or a vendor's data sheet.

Quick Wins

We have quickly worked through the main concepts of using network security tactics to provide access to the myriad of endpoint and mobile devices, so now let's shift to a process to ensure your project succeeds. This is all about progress and avoiding failure, so the best path is to focus your project on establishing an initial quick win, and then gradually build momentum for the technology with expanded deployment.

Step 1: Define Success

We know this seems obvious, but it's amazing how many organizations just start projects without focusing on the problem to solve and how to gauge success. Start every process by making sure everyone is on the same page for what needs to be protected, from what specific threats, via what access mechanisms. You can do a formal threat model or an informal list of use cases. Ultimately to deploy much of anything, you'll need policies that reflect the business needs of the organization. In many cases, a compliance driver is the catalyst for the deployment, so defining how to meet the regulatory mandate needs to be part of the discussion. Regardless of the catalyst, the first step begins by defining and gaining consensus on what success means for this project.

Step 2: Establish Deployment Plan

What's next? Protect the most critical information, of course. In this step, get everyone on the same page regarding where enforcement points will be installed and how you'll phase in the deployment. Understand up front that you will be wrong — what makes the most sense changes as you go through the project. This isn't about carving anything in stone — it's thinking ahead of time about the best way to solve your problem — *before* some vendor puts you on a runaway train.

Note that all this work happens *before* you start engaging with vendors. We advocate a strong plan *before* product evaluation. Things change, but if you don't know what you want ahead of time, the odds are you will *never* accomplish it.

Step 3: Technology Evaluation

Now you get to suffer though any number of dog and pony shows to establish your short list of vendors. We suggest keeping the meetings focused and making sure you do some homework before sitting with a vendor. Then you'll at least know when they are blatantly pulling your leg.

Step 4: PoC

When dealing with complicated technology, we **always** recommend a proof of concept before buying anything. Given the number of integration points for Network Access Control, you'd be crazy not to ensure each vendor can work with your existing gear.

We also believe the PoC needs to be customer driven; which means *you* define the use cases, integration points, and management tasks to be tested — not the vendor. Surprisingly enough, vendors have an unfortunate tendency to guide you toward their strengths. You need to stay focused on solving *your* problem. Be particularly wary of user experience and day-to-day operations, because once you buy something you'll be living with it every day for quite a while. Also pay attention to the documentation generated by the solution, given that inevitably an assessor will show up and you'll need to substantiate the controls you've implemented.

And ensure you have the operational groups on board during the PoC — particularly the network and endpoint folks. Implementing NAC (or something like it) impacts both these areas — often quite significantly. The last thing you need is another group sabotaging your efforts because you didn't line up support early in the process and be sure to document this support, given some folks get a case of forgetfulness when it comes time to formally approve the deployment.

Step 5: Initial Deployment/Quick Win

At this point, after you have selected and bought technology (yes, we skipped a bunch of steps, including actually buying the gear), you need to roll it out. For NAC, we generally recommend an initial focus on visibility. This provides dashboards and reports about what devices are connecting, where they are going, and what they are doing. Gradually enforcement policies for some classes of users/devices can be introduced — once you figure out where the biggest exposures are, based on real usage rather than a theoretical threat model. We favor visibility first because this is a quick win. Breaking users' ability to get onto the network and do work qualifies as a *big loss*.

To take it farther, given the sensitivity around mobile devices, a logical place to start is monitoring them on your network. In our experience this is pretty enlightening, and clearly drives the first set of access control policies. Alternatively you could scrutinize guest access or folks coming in on the VPN from unprotected networks.

We aren't religious about where you start, but make sure you focus on a place where you *know* bad stuff is happening. This way you get proof of the bad stuff and then take quick action to remediate it, which becomes a quick win. Next you can focus on another area of bad stuff and build momentum for the technology and project, all the while generating fancy reports detailing the effectiveness of the new tool.

Summary

We have spent considerable time thinking about the impact of *any* computing (providing access from anywhere, at any time, on any device) on how to protect our networks. These emerging requirements — especially given the avalanche of consumer-oriented mobile devices — are driving us to provide Network Access Control capabilities on our networks. Whether implementing a specific NAC device or using your existing switching and security infrastructure (or a combination of the two), you need to enforce who, what, and how network-based resources are accessed, and guard against unauthorized access to your most critical information.

This involves a wide-ranging project involving risk assessment, sensitive data classification, policy definition (based on likely use cases and threat models) to generate a solution to the business requirements driving widespread access on all of these devices. Once the policies are established, then you'll consider a variety of choices about integrating with the existing network and security infrastructure, as well as endpoint/mobile device management, depending on the level of remediation required on out-of-policy devices.

There are many potential issues regarding this integration and remediation which must be identified and addressed during the procurement process, so focus on a modest initial rollout which proves the concept for network-based access control and builds momentum though quick wins. Keep in mind, you don't have to eat the elephant in one bite, rather a better approach is to phase in the technology by focusing on visibility initially and then implementing controls (limiting or blocking) depending on the sensitivity of the network resources and associated data.

It sounds easy, and on paper it is. You'll find real life a bit more complicated, but as long as you go into the project focused on the business problem you are solving, you can succeed.

About the Analyst

Mike Rothman, Analyst/President

Mike's bold perspectives and irreverent style are invaluable as companies determine effective strategies to grapple with the dynamic security threatscape. Mike specializes in the sexy aspects of security, such as protecting networks and endpoints, security management, and compliance. Mike is one of the most sought-after speakers and commentators in the security business, and brings a deep background in information security. After 20 years in and around security, he's one of the guys who "knows where the bodies are buried" in the space.

Starting his career as a programmer and a networking consultant, Mike joined META Group in 1993 and spearheaded META's initial foray into information security research. Mike left META in 1998 to found SHYM Technology, a pioneer in the PKI software market, and then held executive roles at CipherTrust and TruSecure — providing experience in marketing, business development, and channel operations for both product and services companies.

After getting fed up with vendor life, Mike started Security Incite in 2006 to provide a voice of reason in an over-hyped yet underwhelming security industry. After taking a short detour as Senior VP, Strategy at eIQnetworks to chase shiny objects in security and compliance management, Mike joined Securosis with a rejuvenated cynicism about the state of security and what it takes to survive as a security professional.

Mike published *The Pragmatic CSO* <http://www.pragmaticcso.com/> in 2007 to introduce technically oriented security professionals to the nuances of what is required to be a senior security professional. He also possesses a very expensive engineering degree in Operations Research and Industrial Engineering from Cornell University. His folks are overjoyed that he uses literally zero percent of his education on a daily basis. He can be reached at mrothman (at) securosis (dot) com.

About Securosis

Securosis, L.L.C. is an independent research and analysis firm dedicated to thought leadership, objectivity, and transparency. Our analysts have all held executive level positions and are dedicated to providing high-value, pragmatic advisory services.

Our services include:

- *Primary research publishing*: We currently release the vast majority of our research for free through our blog, and archive it in our Research Library. Most of these research documents can be sponsored for distribution on an annual basis. All published materials and presentations meet our strict objectivity requirements, and follow our Totally Transparent Research policy.
- *Research products and strategic advisory services for end users*: Securosis will be introducing a line of research products and inquiry-based subscription services designed to assist end user organizations in accelerating project and program success. Additional advisory projects are also available, including product selection assistance, technology and architecture strategy, education, security management evaluations, and risk assessments.
- *Retainer services for vendors*: Although we will accept briefings from anyone, some vendors opt for a tighter, ongoing relationship. We offer a number of flexible retainer packages. Example services available as part of a retainer package include market and product analysis and strategy, technology guidance, product evaluations, and merger and acquisition assessments. Even with paid clients, we maintain our strict objectivity and confidentiality requirements. More information on our retainer services (PDF) is available.
- *External speaking and editorial*: Securosis analysts frequently speak at industry events, give online presentations, and write and/or speak for a variety of publications and media.
- *Other expert services*: Securosis analysts are available for other services as well, including Strategic Advisory Days, Strategy Consulting engagements, and Investor Services. These services tend to be customized to meet a client's specific requirements.

Our clients range from stealth startups to some of the best known technology vendors and end users. Clients include large financial institutions, institutional investors, mid-sized enterprises, and major security vendors.

Additionally, Securosis partners with security testing labs to provide unique product evaluations that combine in-depth technical analysis with high-level product, architecture, and market analysis. For more information about Securosis, visit our website: http://securosis.com.