Shining a Light on Shadow Devices

Version 1.6
Released: May 27, 2016

Securosis, L.L.C. 515 E. Carefree Highway 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, but has been enhanced, reviewed, and professionally edited.

Special thanks to Chris Pepper for editing and content support.

This report is licensed by ForeScout Technologies.

www.forescout.com

ForeScout Technologies, Inc. is transforming security through visibility. ForeScout offers Global 2000 enterprises and government organizations the unique ability to see devices, including non-traditional devices, the instant they connect to the network. Equally important, ForeScout lets you control these devices and orchestrate information sharing and operation among disparate security tools to accelerate incident response. Unlike traditional security alternatives, ForeScout achieves this without requiring software agents or previous device knowledge. The company's solutions integrate with leading network, security, mobility and IT management products to overcome security silos, automate workflows and enable significant cost savings.

Copyright

Shining a Light on Shadow Devices

Table of Contents

The Exponentially Expanding Attack Surface

It seems like a long time ago, but do you remember when Internet-connected devices were all desktop PCs? Then we got fancy and started issuing laptops. That's what connected to our networks. Life was simpler then. But you don't have time for nostalgia. You are too busy getting a handle on the explosion of devices connected to your networks now, accessing your data.

Here is just a smattering of what we see:

- **Mobile Devices:** Supporting smartphones and tablets is old news, mostly because we cannot remember a time when they weren't on our networks. But despite their short history, the impact of smart devices on mobile networking and security cannot be understated. Even more challenging, these devices can connect directly to the cellular data network for a path around your security controls.
- **BYOD:** Then someone decided it would be cheaper to have employees use their own devices, and Bring Your Own Device (BYOD) became a thing. You can have employees sign paperwork giving you permission to control their devices and install software, but in practice they get very cranky (understandably) whenever you prevent them from doing something on *their* own devices. Balancing protection of corporate data against antagonizing employees has been challenging.

You don't have time for nostalgia. You are too busy getting a handle on the explosion of devices connected to your networks now, accessing your data.

- **Other Office Devices:** Printers and scanners have been networked for years. But as more sophisticated imaging devices emerged we realized their on-board computers and storage were insecure. They became targets, providing attackers with beachheads on your networks.
- **Physical Security Devices:** The new generation of physical security devices (cameras, access card readers, etc.) is largely network connected. It's great that you can grant access to a locked-out employee from your iPhone on a golf course, but much less fun when attackers grant *themselves* access via a security device using the default password.

- **Control Systems and Manufacturing Equipment:** The connected revolution has made its way to shop floors and facilities areas as well. It could be a sensor collecting information from factory robots or a warehouse system, but many of these devices are networked too — so they can be attacked. You may have heard of StuxNet targeting centrifuge control systems. That's the kind of thing we're talking about.
- **Healthcare Devices:** If you go into any healthcare facility nowadays, monitoring devices and even some treatment devices are managed through network connections. There are jokes to be made about taking over shop floor robots, but who really cares, aside from screenwriters? But the ramifications of attacks on medical devices are far more severe.
- **Connected Home:** Whether it's a thermostat, security system, or home automation platform, the modern expectation is that you will manage it from wherever you are. That means a network connection and access to the Intertubes. What could *possibly* go wrong? But don't worry — you'll find out!
- **Cars:** An automobile can now use either your smartphone connection or its own cellular link to connect to the Internet for traffic alerts, music, news, and other services. It can transmit diagnostics as well. This is pretty cool, but recent stunt hacking has proven that moving automobiles can also be attacked and controlled remotely. Uh-oh.

There will be billions of devices connected to the Internet over the next few years. They all present attack surface on your technology infrastructure. And you cannot fully know what is exploitable in your environment, because you don't know about these devices living in the 'shadows'.

The industry wants to dump all these devices into a generic Internet of Things (IoT) bucket, because IoT is the buzzword *du jour*. All these network thermostats and washing machines have become the latest Chicken Little, poised to bring down the sky. It turns out the sky already fell — networks are *already* too vast to fully protect. The problem is getting worse by the day, as pretty much anything with a chip in it gets networked, and the rest gets chipped. So instead of a manageable environment, you need to protect Everything Internet.

All these network thermostats and washing machines have become the latest Chicken Little, poised to bring down the sky. It turns out the sky already fell — networks are already too vast to fully protect.

Anything with a network address can be attacked. Fortunately better fundamental architectures (especially for mobile devices) make it harder to compromise new devices than traditional PCs (whew!), but sophisticated attackers don't seem to have trouble compromising any device they can reach. And that says nothing of all the devices whose vendors have paid little or no attention to security to date. Healthcare and control system vendors, we're looking at you! These devices have porous defenses, if any, and once an attacker gains presence on the network, they have a bridgehead to work their way to their ultimate target.

In the Shadows

So what? You don't even *have* medical devices or control systems, so why should you care? The sad fact is that what you don't see *can* hurt you. Your entire security program has been built to protect what you can see with traditional discovery and scanning technologies. The industry has maintained a very limited concept of what you should be looking for — largely because that's all security scanners could see. The current state of affairs is that you run scans every so often and see new devices emerge. You test them for configuration issues and vulnerabilities, and then add those issues to the end of an endless list of things you'll never have time to finish.

Unfortunately visible devices are only *some* of the network-connected devices in your environment. There are hundreds, and quite possibly thousands, of other devices you don't know about on your network. You don't scan them periodically, and you have no idea of their security posture. Each one can be attacked, and might provide an adversary with opportunity to gain a presence in your environment. Your attack surface is much larger than you thought.

These shadow devices are infrequently discussed and rarely factored into discovery and protection programs. It's another Don't Ask, Don't Tell approach, which never seems to work out well.

We haven't yet published anything on IoT devices (or Everything Internet), but it is time. Not because we currently see many attacks in the wild. But most organizations we talk to have not adequately prepared for attacks, so they will scramble… as usual. We have espoused a visibility-then-control approach to security for over a decade. Now it's time to get a handle on the visibility of all devices on your network, so when you need to, you will know what you have to control — and how.

Attacks

What is the real risk of shadow devices connecting to your networks? It's clear most organizations don't take these risks seriously. They certainly don't have workarounds in place in case these devices are compromised, or proactively segment environments to ensure that compromising them doesn't provide opportunity for attackers to gain presence and move laterally through their environments to critical data. We will dig into three broad device categories to understand what the attacks look like.

Peripherals

Do you remember how cool it was when the office printer got a WiFi connection? Suddenly you could put it wherever you wanted, preserving the Feng Shui of your office, instead of having it tethered to a network drop. And when printer makers started calling their products *image servers*, not just printers? That was when they started becoming more intelligent — and more tempting targets.

> But what is the risk of taking over a printer? It turns out that even in the paperless offices of the future, organizations still print sensitive stuff, and things they don't want to keep might be scanned for archival.

But what is the risk of taking over a printer? It turns out that even in the paperless offices of the future, organizations still print sensitive stuff, and things they don't want to keep might be scanned for archival. Whether going in or out, sensitive content is hitting image servers. Many of them store the documents they print and scan until their memory (or embedded hard drive) is overwritten. Sensitive documents persist on devices, accessible to anyone with access to the device, whether physical or remote.

Even better, printers are vulnerable to common wireless attacks like the evil twin, where a fake device with a stronger wireless signal impersonates the real printer. So devices connect (and print) documents to the evil twin instead of the real printer (or, for extra sneakiness, *both*). The same attack works with routers too, but there the risk is even broader. Nice. But that's not all! The devices typically use some kind of stripped-down UNIX core, and many organizations don't change the default passwords on their image servers, enabling attackers to trigger remote firmware updates and install compromised versions of the printer OS. For extra fun, these imaging devices now connect to cloud-based services to email documents, so they have all the plumbing to act as spam relays.

Most printers use similar open source technologies to provide connectivity, so generic attacks work against a variety of manufacturers' devices. These peripherals can be used to steal content, attack other devices, and provide a foothold inside your network perimeter. That makes them both *direct* targets with information interesting to attackers, and *indirect* targets which can be used to move laterally within an organization.

Unfortunately these attacks aren't just theoretical. We have seen printers hijacked to spread inflammatory propaganda on college campuses, and Chris Vickery showed a proof of concept to access a printer's hard drive remotely.

Our last question is what kind of security controls run on imaging servers? The answer is: none really. To be fair, vendors have started looking at protecting peripherals more seriously, and the impacted vendors have been reasonably responsive in patching the attacks mentioned above. But these products do not get the same scrutiny as PC devices, or even other connected devices we will discuss below. Imaging servers get minimal, if any, security assessment before coming to market.

We aren't just picking on printers. Pretty much every intelligent peripheral is similarly vulnerable, because they all have operating systems and network stacks which can be attacked. It's just that offices tend to have dozens of printers, which are frequently overlooked during risk assessment.

Medical Devices

If printers and other peripherals seem like low-value targets, let's discuss something a bit higher-value: medical devices. In our era of increasingly connected medical devices — including monitors, pumps, pacemakers, and pretty much everything else — there hasn't been much focus on product security, except in a few cases where external pressure was applied by regulators. These devices either have TCP/IP network stacks or can be configured via Bluetooth — neither of which is particularly well protected.

Despite the very real risk of medical devices being targeted to attack patients, we believe most of the current risk involves information. User data is much easier for attackers to monetize: medical profiles have a much longer shelf-life and much higher value than most other financial information.

The most disturbing attacks threaten patient health. There are all too many examples of security researchers compromising infusion and insulin pumps, 'jackpotting' drug dispensaries, and even the legendary Barnaby Jack messing with a pacemaker. We know one large medical facility that took it upon itself to hack all its devices and deliver a list of issues to manufacturers. But there has been no public disclosure of results, or whether device manufacturers made changes to make their devices safer.

Despite the very real risk of medical devices being targeted to attack patients, we believe most of the current risk involves information. User data is much easier for attackers to monetize: medical profiles have a much longer shelf-life and much higher value than most other financial information. So ensuring that Protected Health Information is adequately protected remains a key concern.

That means making sure there aren't any leakages in these devices, which is not easy, even when running a full penetration test. On the positive front, many of these devices have purpose-built operating systems, so they cannot easily be used as pivot points for lateral movement within the network. Yet few have any embedded security controls to ensure data does not leak. Further complicating matters, those not using a purpose-built OS typically use deprecated operating systems such as Windows XP or even Windows 2000 (yes, seriously), and outdated compliance mandates often mean they *cannot* be patched without recertification. So administrators often cannot update these devices without breaking the law, and are forced to hope for the best. We can all agree that hope isn't a sufficient strategy.

With lives at stake, medical device makers are starting to talk about proactive security testing. Similarly to the way a major SaaS breach could seriously threaten the larger SaaS market, medical device makers should understand what is at risk, especially in terms of liability. But that doesn't mean they understand how to solve the problem. So the burden lands on customers (hospitals and doctors) to manage their medical device inventories, and ensure they are not misused to steal data or harm patients.

Industrial Control Systems

The last category of shadow devices we will consider is control systems. These devices include SCADA systems running power grids, warehousing systems ensuring the right merchandise is picked and shipped, manufacturing systems running robotics, and heavy building machinery. All these devices are networked (whether directly or indirectly) in today's advanced factories, so there is attack surface to monitor and protect.

We know these systems can be attacked. Stuxnet was a very advanced attack on nuclear centrifuges. Once within the nuclear facility's network, adversaries compromised a number of different types of control systems to access centrifuges and break them. In a recent attack on a German blast furnace, control systems were compromised and general failsafes were inoperable: the facility went offline while they cleaned the systems up, impacting product delivery.

In both cases, and likely many others that haven't been publicized, the adversaries were very advanced. They need to be — to attack a centrifuge like Stuxnet you need your own centrifuges to test on, and those aren't easy to find on eBay. You cannot just load a blast furnace into the pick-up Saturday morning for a penetration test.

That may comfort some people, but it shouldn't. The fact is that control system defenders aren't dealing with the Metasploit crowd, but trying to repel well-funded and capable adversaries. These organizations need a very clear idea of what their attack surface looks like and some way of monitoring their devices — they cannot rely on compliance mandates to require or blueprint advanced security for their systems.

Another consideration for control systems is their brittle nature. They are hard to test because you could bring down the system while trying to figure out whether it’s vulnerable.

Another consideration for control systems is their brittle nature. They are hard to test because you could bring down the system while trying to figure out whether it’s vulnerable. Most organizations don’t like that trade-off, so they don’t test directly. This means you need indirect techniques — definitely to figure out how vulnerable your systems are, and probably to discover and monitor them as well.

Seeing into the Shadows

An increasing number of the devices connecting to your networks aren't traditional computers, so they cannot be scanned and assessed for security issues the same way. These *shadow devices* need to be taken into consideration when designing security architectures.

We have explained how peripherals, medical devices, and control systems can be attacked. We showed that although traditional malware attacks on traditional computing and mobile get most of the attention in IT security circles, other devices shouldn't be ignored. As with most things, it's not a question of *if* but *when* lower-profile devices will be used to perpetrate a major attack.

How can you figure out your real attack surface? You need to go back to Security 101, which still starts with visibility, and moving on to control once you figure out what you have and how it is exposed.

So now what? How can you figure out your real attack surface, and then move to protect the systems and devices providing access to your critical data? You need to go back to Security 101, which still starts with visibility, and moving on to control once you figure out what you have and how it is exposed.

Risk Profiling

Your first step is to shine a light into the 'shadows' on your network to gain a picture of all devices and develop a 'profile' of the risk each device poses to your organization. You have a couple ways to gain this visibility:

1. **Active Scanning:** You can run a scan across your entire IP address space to find out what's there. This can be a serious task for a large address space, consuming resources while you run your scanner(s). This process only happens periodically — as a rule scanners do not run continuously on internal networks. Keep in mind that some devices, especially ancient control systems, were not built with resilience in mind, so even a simple vulnerability scan can knock them over.

2. **Passive Monitoring:** The other alternative is to watch for new devices by monitoring network traffic. This assumes that you have access to all traffic on all networks in your environment, and that new devices will communicate to *something*. Pitfalls of this approach include needing access to the entire network, and the possibility of new devices spoofing other network devices to evade detection. On the plus side, you won't knock anything over

by listening and if the attack hasn't been detected, once the malware starts using the network to communicate, you can see it.

But how to gain full visibility into *all* devices on the network is not an either/or question. There is a time and place for active scanning, but care must be taken to not take brittle systems down or consume undue network resources. We have also seen many scenarios where passive monitoring is needed to find new devices quickly once they show up on the network.

Once you have full visibility the next step is to identify devices. You can certainly look for device type indicators during active scanning. This is harder when passively scanning, but devices can often be identified by traffic patterns and other packet indicators. A critical selection criteria for passive monitoring technology is the vendor's ability to identify the bulk of devices likely to show up on your network. Obviously in a very dynamic environment some fraction of devices cannot be identified through scanning or monitoring network traffic. But you want these devices to be a *small* minority because anything you fail to identify through scanning requires manual intervention.

Once you know what kind of device you are dealing with, you need to start evaluating risk — a combination of the device's *vulnerability* and *exploitability*. Vulnerability is what could possibly happen. An attacker can do certain things with a printer which are impossible with an infusion pump, and *vice-versa*, so device type is key context. You also need to assess security vulnerabilities within the device. They may warrant an active scan for more detail. As we warned above, be careful with active scanning to avoid compromising device availability. Keep in mind that passive scanning requires quite a bit more interpretation, and is subject to higher false positive rates.

Exploitability depends on the security controls and/or configurations already in place. A warehouse picker robot might run embedded Windows XP, which is a security train wreck. But if the robot also uses an application whitelist malicious code cannot execute, so vulnerabilities might not be exploitable. The other key to exploitability is *attack path*. If an external entity cannot access the warehouse system because it has no Internet-facing networks, even a vulnerable picker robot poses relatively little risk unless the physical location is attacked. But keep in mind that WiFi footprints often extend outside the building.

> Determining risk is all about prioritization. You only have so many resources, so you need to choose which issues to fix wisely, and evaluating risk is the best way to allocate scarce resources.

The final aspect of determining risk to a device is what it can access. If a device has no access to anything sensitive, again it poses little risk. Of course that assumes your networks are sufficiently isolated. Determining risk is all about prioritization. You only have so many resources, so you need to choose which issues to fix wisely, and evaluating risk is the best way to allocate scarce resources.

Controls

Once you know what's out there in the shadows, your next step is to figure out whether, and perhaps how, to protect it. This again comes back to the risk profiles discussed above. It doesn't make much sense to spend a lot of time and money protecting devices which don't present much risk to the organization. But in case a device *does* pose sufficient risk, how will you go about protecting it?

First things first: you should be making sure each device is configured in the most secure fashion. That sounds trivial but we need to mention it anyway, because it's shocking how many devices are exploited thanks to open services that can easily be turned off. Once you have basic device hygiene taken care of, here are some more ways to protect a device:

1. **Active Controls:** The first and most direct way to protect a shadow device is by implementing an active control on it. The available controls vary depending on device type and embedded operating system. The most reliable protection is to prevent execution of unauthorized programs: whitelisting. Commercial products are available for embedded Windows devices, and a few options exist (both commercial and open source) for other operating environments. There are also other defenses, including malware detection and even some forensic offerings, to determine what's *really* happening on devices. Just keep in mind that active controls consume resources and can impair device stability. So you'll want to exhaustively test any controls to be implemented on these devices to ensure you understand their impact.

2. **Network Isolation:** These devices are on the network, which means traditional network security defenses can (and should) be used in conjunction with active controls. This involves using firewalls and/or routers & switches to provide a measure of isolation between device networks and data stores.

3. **Egress Filtering:** We also recommend egress filtering, and the inevitable presence of shadow devices is yet another reason. You can detect command and control traffic, as well as remote access, by looking at outbound network traffic. Remember, these devices aren't inherently malicious, so if something bad it happening it's because an adversary is controlling the device, which means it's communicating to a bot network or other controller outside your network — and that can be tracked.

An increasing limitation for network-oriented controls is encrypted traffic. We have published research on dealing with encrypted networks — ways to peek into encrypted traffic streams. When deciding between network and active controls for devices, take encryption into account — network-based controls are easier because they don't require device interaction, but may introduce blind spots, particularly with encryption in use.

Automation

Finally we should mention the challenges of implementing fixes and workarounds — not only when dealing with shadow devices, but for all networked devices now. There just isn't enough staff to address all requirements, which means organizations need to think creatively about how to make limited staff more productive. One effective way to amplify staff effort is automation. Existing controls can be reconfigured based on specific rules. Of course automation of existing technologies requires integration with existing security controls, but lately we have seen considerable innovation in this area, born of necessity.

There just isn't enough staff to address all requirements, which means organizations need to think creatively about how to make limited staff more productive. One effective way to amplify staff effort is automation.

Our last step in figuring out your strategy to gain visibility and control over shadow devices is to determine which fixes can be automated, and establish a strategy for that automation. We caution again that automation is a double-edged sword, which must be used carefully to ensure the 'cure' isn't much worse than the symptoms. So ensure sufficient testing, and have reasonable expectations for what you can automate — in both short and long terms.

Summary

As we explained, we will see an explosion of devices connecting to the networks you are tasked to protect. Many are not built to be scanned or protected by traditional security technologies. These *shadow devices* are an under appreciated attack vector now, given the seemingly very urgent malware and traditional attacks most organizations see daily. But adversaries are not religious about how they steal information or gain a foothold in your organization. They'll do whatever it takes to achieve their mission.

Without sufficient visibility of these non-traditional devices you will be blind to potential attacks because you can't protect what you cannot see. So we recommend you start by focusing on *visibility* to discover the entirety of your attack surface, and then determine the best means of providing controls to protect your environment.

If you have any questions on this topic, or want to discuss your situation specifically, feel free to send us a note at info@securosis.com.

About the Analyst

Mike Rothman, Analyst and 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 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. 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, LLC 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 conform to 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 assessment.

- **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. Services available as part of a retainer package include market and product analysis and strategy, technology guidance, product evaluation, and merger and acquisition assessment. 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 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 tend to be customized to meet a client's particular 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/>.