armis

SECURING UNMANAGED DEVICES IN THE ENTERPRISE

EXECUTIVE OVERVIEW

Modern enterprises are increasingly inundated, and simultaneously reliant upon, devices that don't fit into the traditional model of management and security. Devices of all types are increasingly network-capable by default. And in many cases, contain a blend of wired and wireless connectivity options. Yet you can't install a traditional security agent on these devices, and patching or updating their operating systems can be extremely challenging for a variety of reasons. Taken together, this presents a near perfect storm of risk — devices that are accessible, vulnerable, and unprotected.

And while these devices do not fit our traditional models, they are not rare. IoT and smart devices like personal end-user devices, laptops, smartphones, or office mainstays such as printers, Smart TVs, and even the very switches and server controllers at the heart of our networks, all fit the definition of an "unmanaged device." These devices are essential to the business and can't be ignored or simply denied access. This reality has increasingly put security teams in an untenable position. Without the proper tools, security teams are faced with time-consuming work, policies that can be disruptive to the business, and ultimately are not effective at stopping threats.

Instead we need to recognize that the evolution of our devices demands a corresponding evolution in our security architecture. If security teams are to empower the business, they likewise must be empowered with a security architecture that confronts the realities of modern devices and networks.

In this paper we will take a deeper look at the challenges and realities of unmanaged devices, and propose a new approach that allows security teams to proactively secure the environment, while safely enabling the business to use the devices that it needs.

GETTING TO KNOW THE INVISIBLE LEGION

Unmanaged devices are largely defined by what they don't do. Put in the most reductive terms, unmanaged devices don't run a traditional security agent. But this definition casts a wide net that covers an incredibly diverse set of devices ranging from the simplest of devices in our networks with tiny amounts of memory to some of the most sophisticated devices running custom operating systems. The list below provides a high-level overview of some of the most common unmanaged devices found in enterprises:

- **Office Devices and Peripherals** - Printers, VoIP phones, TV screens and monitors, Bluetooth keyboards, headsets, etc.
- **IoT Devices** - HVAC systems, security systems, lighting systems, cameras, refrigerators, vending machines, etc. This is itself a very broad category, but we will include devices that are networked primarily for purposes of collecting telemetry data and remote management.
- **Personal or Consumer Devices** - Smart devices, gaming consoles, streaming media, or digital assistants (Apple TV, Slingbox, Amazon Echo, etc).
- **Industry-Specific Devices** - Industrial Control Systems, medical devices (patient monitoring systems, mobile imaging systems, infusion pumps, communication badges, etc), retail (barcode scanners, POS system, loss prevention, etc).
- **IT Infrastructure** - Routers, switches, firewalls, baseboard management controllers of servers. While these devices are not typically thought of as unmanaged devices, they do not support traditional security agents and have been increasingly targeted by sophisticated attackers.

The danger is that we can easily think of each example as a separate, unrelated corner case. And this is simple human nature. An employee's personal Android device, a wireless security camera, and a network switch are such different devices, we rarely think of what they share in common. Regardless of the fact they all fall outside of our traditional security models for different reasons, collectively they account for a surprisingly large part of the devices on our networks. While these devices may not be "standard", they certainly are not rare. So instead of focusing on how these devices are different, let's briefly consider the things that they share in common

No Security Agent

First, as we mentioned above, we have to consider any device on which you can't install a traditional security agent. In most cases this is due to one or more of the following reasons:

- **Non-standard operating system** - By non-standard we mean any operating system not supported by endpoint security products. Depending on the endpoint security vendor this can mean anything not running Microsoft Windows or Apple's OS X. Support for mobile operating systems vary widely as does support for the many Linux distributions. Additionally, devices such switches, routers, firewalls and industrial control systems can run custom operating systems.
- **Limited resources on the device** - Many devices simply lack the memory or processing needed to run an agent. Many IoT devices exist as "systems on a chip," and are designed with the bare minimum of resources needed to perform their task. And while you can't install a security agent on these devices, they are highly valuable targets to attackers.
- **Devices not owned by the enterprise** - This can include devices owned by employees, contractors, consultants, or vendors. Many organizations often have devices that are shared between partners or subsidiaries.

Unpatchable Devices

In addition to not running a security agent, many devices are impossible, or at best, highly impractical to patch. As we saw in the previous section, this can include devices that are not directly owned or under the control of the enterprise, such as devices shared between organizations. In other cases, legacy devices may be unpatchable simply due to age or lack of support. The longevity of Windows XP and protocols such as NTLM within the enterprise is a good reminder of just how long certain devices can survive past their expiration date.

However sometimes the bigger issue is that patches are often not available in the first place. Unlike standard operating systems where updates are frequent and often automated, patches for unmanaged devices tend to be very few and far between. IoT devices in particular are notoriously slow to deliver patches, and are often delivered only in response to a public attack against the device. It is not at all uncommon for devices to go years between patches and take months to respond even once a vulnerability is disclosed. The end result are devices with extremely antiquated operating systems and modules with well-known documented vulnerabilities.

UNMANAGED DEVICES AND THE CORROSION OF THE TRADITIONAL SECURITY MODEL

While unmanaged devices may be surprisingly pervasive in the enterprise, this by itself is not necessarily cause for alarm. What is of concern are the ways that these devices have undermined industry-standard approaches to security. And when an architecture begins to fail, security teams are often left in a particularly precarious spot. If the architectural issues aren't addressed, teams are left constantly trying to plug leaks in a failing dam. This often leads to time-consuming, expensive, and ultimately ineffective efforts. In this section we will take a look at how unmanaged devices have created gaps in the traditional security architecture and how attackers are taking advantage of them.

An Easily Discoverable Attack Surface

We need to realize that just because these devices may be beyond the view of traditional management tools, they are not invisible to attackers. Search engines such as Shodan make it trivial for anyone to find virtually any type IoT or connected device. This means that if/when an attacker finds a vulnerability in a wireless camera, for example, he can quickly find vast numbers of suitable, real-world targets. This is not a problem caused by Shodan or similar sites. They are simply search engines scanning the Internet, doing what any search engine can do. The issue is not that Shodan exists, but rather that any attacker could do the same thing.

Not only are these devices visible to hackers, but they are active targets for attacks. This was the case with the Mirai attack. Recent reports show that 46% of breaches or security incidents were associated with IoT security.[1] Additional data reveals a 280% increase in attacks focused on IoT devices in the first half of 2017.[2]

Perimeter Evasion

While unmanaged devices are almost by definition beyond the reach of endpoint security, they also create problems for network security. Network security has historically relied on the assumption that traffic can be funneled to logical choke points where the traffic can be analyzed and policy enforced. The network perimeter is the classic example where all traffic is funneled to a few ingress/egress points.

However, when it comes to unmanaged devices, Wi-Fi connectivity is often the default network option. And this means that unmanaged devices are constantly and directly exposed to untrusted devices and networks. Wireless has the uncomfortable trait of completely ignoring our carefully constructed wired perimeters and network segmentation. And while wireless is certainly nothing new to the enterprise, we have to remember that unmanaged devices are uniquely unprotected by endpoint security. To an unmanaged device, endpoint security is not compatible, and network security can be avoided. It is this combination that makes unmanaged devices so antithetical to traditional security architectures.

Modeling the Threats

The exposure of unmanaged devices opens the organization to both inbound and outbound risks. Information can potentially leak from the enterprise out without ever touching an organization's firewall, IPS, or DLP solution. One of the most common responses when considering the security of unmanaged devices is that the devices themselves don't have any real data on them or privileges within the network. The counterpoint is that we have to remember that any device with a network stack can be turned into a proxy. At an absolute minimum, a wirelessly networked device can provide a relay into and out of a network that is virtually invisible to traditional security. It should give us all pause when our best-case scenario is an invisible command-and-control or exfiltration channel.

Likewise, wireless connectivity opens the door to a wide variety of additional malicious techniques that normally would require an attacker to have a presence on the network. ARP poisoning, man-in-the-middle techniques are instantly fair game.

Unfortunately, the risks of unmanaged devices have gone well beyond the theoretical. Disclosures from WikiLeaks and Shadow Brokers have shown a heavy focus on IoT and unmanaged devices. The WikiLeaks Vault 7 disclosures revealed a variety tools focused on smartphones, smart TVs, and network routers. The Shadow Brokers dumps detailed a variety of exploits and tools aimed at compromising the very firewalls, switches, and infrastructure at the heart of our networks. And regardless of we may feel about these disclosures, attacks like WannaCry illustrate just how disclosed exploits and techniques can be quickly adopted by anyone in the wild.

Unmanaged devices are constantly and directly exposed to untrusted devices and networks.

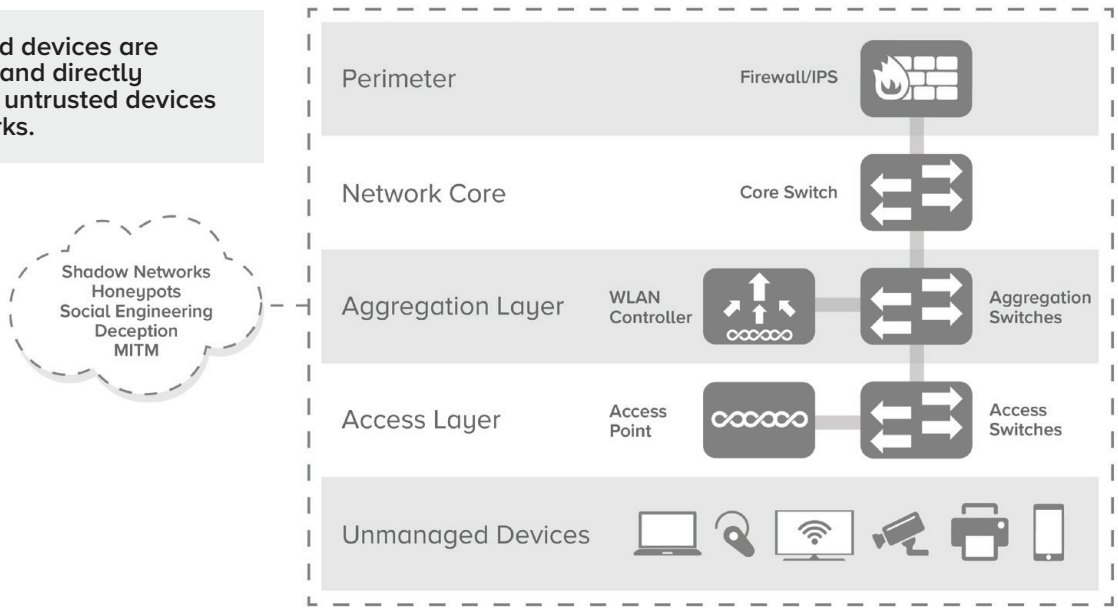

Gaps in Threat Prevention

Given the threats that unmanaged devices face, it is obviously important to monitor them for a wide variety of malicious techniques and behaviors. Is it under attack from an adversary? Does the device show signs of compromise? But much like firewalls, IDS/IPS solutions are often deployed at the network perimeter. And even when they are deployed in internal segments, they are often deployed specifically in front of key applications and do not provide coverage for an entire access layer.

Bringing IPS deeper into the network is often impractical for most organizations. Large numbers of network segments can make it difficult to place sensors so that all traffic is analyzed. The large volumes of internal traffic and the potential exposure to old vulnerabilities would require IPS sensors with high traffic capacities and large signature packs. For most organizations, such an approach is both operationally and cost prohibitive.

Access Control: When the First Line of Defense Becomes the Last

Absent the ability to control threats directly, most organizations have reverted to access control such as 802.1x and NAC as the best option available for securing unmanaged devices. And while these are valid and important technologies, we have to remember that they are part of a network architecture, not a replacement for one. We would never consider protecting our laptops and servers exclusively with pre-access controls, so why would we think it should suffice for unmanaged devices?

Just as is the case with laptops and servers, the majority of the security lifecycle happens after access is granted. If a device is compromised, it will almost assuredly still be able to authenticate to the network. The issue becomes not whether access should be granted, but whether security can detect and stop any suspicious or malicious behavior.

Furthermore, security that focuses exclusively on isolation tends to be painful both for security staff as well as the enterprise. As we have seen, many unmanaged devices are critical to the enterprise, and they need to be connected in order to do their job. Strict isolation policies often limit these devices and ultimately puts security at odds with the business.

Additionally, access control and isolation strategies can quickly become unwieldy in practice and hard to maintain. Controlling unmanaged devices through segmentation and subnets can quickly turn into an exercise in herding cats. There are many devices, and it is easy to miss some along the way. Furthermore, many unmanaged devices will lack the ability to carry an appropriate certificate, forcing staff to revert to whitelisting.

In short order, security staff is left managing a growing list of whitelists and exceptions. In the end, this approach rarely makes anyone happy. Security teams are stuck with kludgy controls, which don't adequately stop threats, and are confrontational with the business. But this situation is the result of compromises being made for an outdated security architecture. So let's take a look at how we can modernize security in a way that addresses the realities of unmanaged devices, and allows security to safely enable the business.

A Loss of Control Where the Action Is

If we look at all of this collectively, we start to see an unsettling trend. The numbers of unmanaged devices are exploding, and attackers are increasingly using them to infiltrate, spread, persist, and ultimately steal. However, the closer we get to this layer of unmanaged devices, our visibility, control, and threat prevention typically gets progressively worse.

At the perimeter, we see all traffic, control what gets in and out, and apply multiple types of threat detection and prevention. Once we move inside the network, this visibility and control drops markedly. Some organizations may have solutions to capture and analyze traffic traversing the core switch, although most of these solutions are deployed as passive detection technologies that lack the ability to actively protect.

The loss of visibility and control accelerates the deeper we go into the network. Visibility into actual traffic drops away almost completely, and we are left mining logs for insight into internal network behavior. By the time we reach our unmanaged devices, the best-case scenario is a NAC solution where most unmanaged devices are simply whitelisted and ignored.

Security visibility and control is the least where unmanaged device density is the greatest.

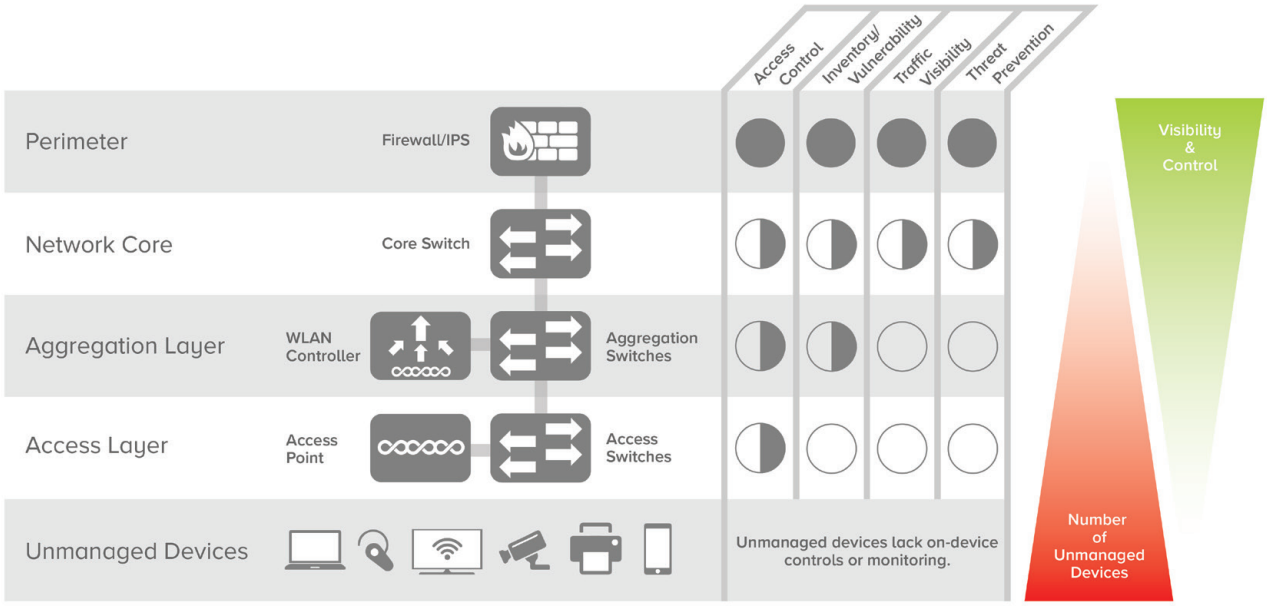

MODERNIZING THE SECURITY ARCHITECTURE FOR ALL DEVICES

Having established some of the challenges related to securing unmanaged devices, we can now turn our attention to the all-important task of designing a solution. This will require us to define high-level requirements for what a solution must do, as well as the more practical side of how those requirements can be met.

To do so, we need to understand not just the realities of our network, but also the realities of the devices we are trying to protect. Most unmanaged devices need to be plug-and-play and are thus designed to make connections easily and automatically. They also exist in a sea of potential connections. Neighboring wireless networks, Bluetooth devices, direct device connections, and personal hotspots just scratch the surface.

Unmanaged devices exist in a sea of potential connections.

Furthermore, device will often beacon for previous connection such as a phone that previously connected to a network at a coffee shop or hotel. Attackers can easily see these beacons and create fake networks to lure devices into connecting automatically. Attackers can be more aggressive and use simple, readily available tools to knock devices off of existing connections so that the attacker can push them to a malicious network.

The good news is that most unmanaged devices have a relatively narrow set of behaviors that they should be allowed. They should have a limited set of common connections, and a similarly limited set of services and actions available to them.

Key Functional Requirements

Ultimately the challenge is to bring a similar level of visibility and control to our unmanaged devices, and safely enable their use without blindly whitelisting them. If we think back to the gaps in our security controls, it quickly becomes obvious that we will need to find the devices that are not secured, appropriately control their connections, monitor their traffic and behavior, and block any malicious behavior. All of this needs to be done in the context of how a device is used so that we can sanction only the actions and privileges that are required to support the business.

Automated Discovery of All Unmanaged Devices

As we have seen, unmanaged devices come in many forms, and security staff are often blind to the presence of these devices. Needless to say, it is hard to secure what you can't see. So establishing visibility is a critical first step. Given that unmanaged devices can be transient such as a device introduced by an employee or contractor, it is important that the device discovery process is both continuous and automated.

Furthermore, we may need to include multiple types of wireless visibility to include both Wi-Fi as well as Bluetooth, and other protocols if necessary. Given that wireless exposes our devices to untrusted devices or networks, we will need

to monitor the wireless environment. We can not assume that all connections will be to the corporate network. For example, a device may be both wired and wirelessly connected. An attacker could use the wireless connection as a relay out of the network that would be invisible if external wireless environment is not monitored.

Profile Device Behavior

Once we can see a device, we next must understand what it does. This means observing device behavior over time to establish baselines, and comparing observed behavior to other devices of a similar type or functional role. This profiling should include an understanding of common network connections, protocols in use and other typical behaviors. This phase is also critical for understanding how a device is used in the enterprise, so that we can establish appropriate policies that truly enable the device and the business.

Establish and Enforce Appropriate Behaviors

Next, we need to proactively control the attack surface presented by unmanaged devices. This will require the organization to establish sanctioned behaviors based on the type of device and its role. At a high level, we need to set what is allowed, and deny the rest.

In many cases, policies can afford for policies to be relatively strict. For example, we would expect a wireless IoT camera to behave far more predictably compared to a human user, and thus it would make sense for our policies to limit the camera to a far more narrow set of behaviors.

Unmanaged devices require us to plug the visibility and control gaps in the security architecture.

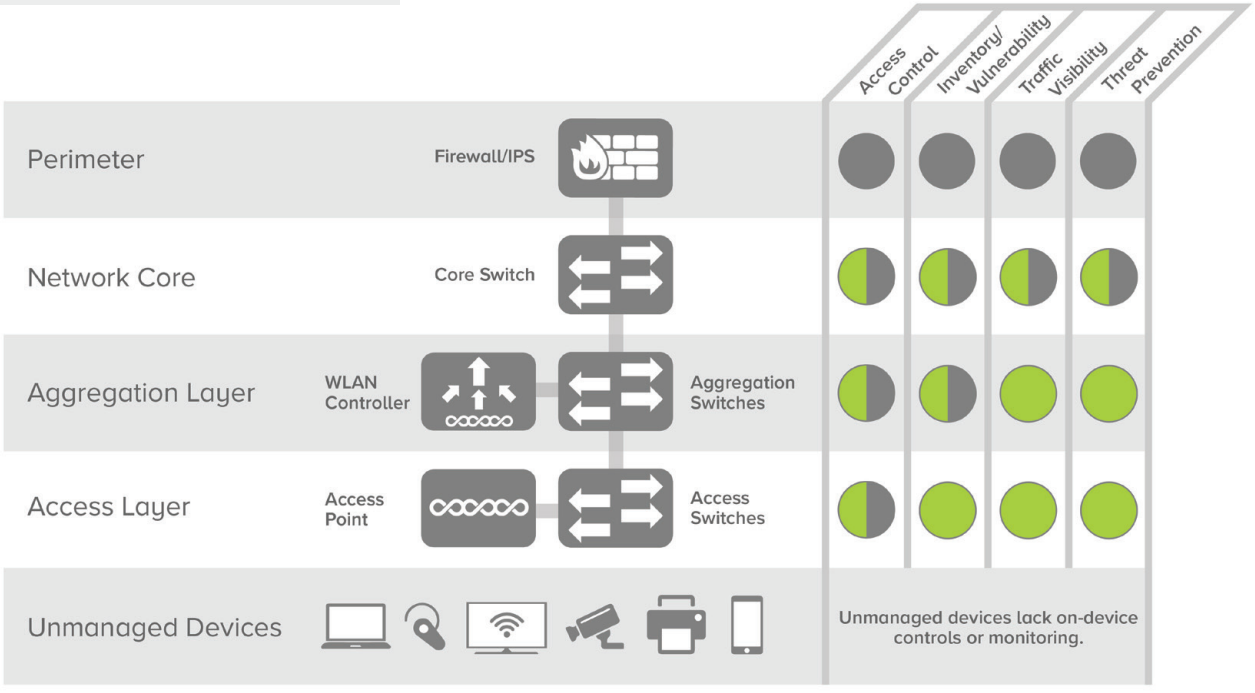

Identify and Stop Malicious Behavior

Having identified and enforced approved behavior, we next need to identify malicious behavior. If our previous step was the analog of a firewall, this phase is the analog of the IPS. Here we should identify the signs of malicious tools and techniques, and also monitor for signs that a device may be compromised. Once a threat is identified, we must have the ability to block the threat automatically. If a device is acting as exfiltration channel of the network, we obviously need to stop the flow of data automatically in order to mitigate damage.

Automated Actions While Minimizing Manual Investigation

Enterprises typically have several times more active IP addresses than users, and this provides a basic indicator of the extent to which devices outnumber human end users. The scale of devices being monitored means that any solution must provide definitive detections and automated enforcement. Any solutions that routinely depend on staff to manually investigate anomalies is very likely to overwhelm staff.

Architectural Fit

It is important to remember that our goal is to modernize our security architecture, and not to "rip and replace" it. As a result, we need to consider how any new solutions will integrate and work with the existing security architecture.

Monitoring From the Wireless Access Layer Up

In order to ensure visibility of device connections coming into our network as well as those connecting to outside devices, we will need to monitor the wireless environment directly. This means that we will need to create visibility that extends from the access layer up, and not form the perimeter down.

This visibility can be established in a variety of ways such as through dedicated wireless sensors, integration with with the existing wireless infrastructure, and even integration with endpoint devices. Each of these approaches can have certain advantages and disadvantages, which are beyond the scope of this paper. However, the key is that in order to see device behavior and to detect abnormal or malicious techniques, a solution will need direct visibility into the environment being protected.

Leverage the existing infrastructure to gain visibility and control over unmanaged devices.

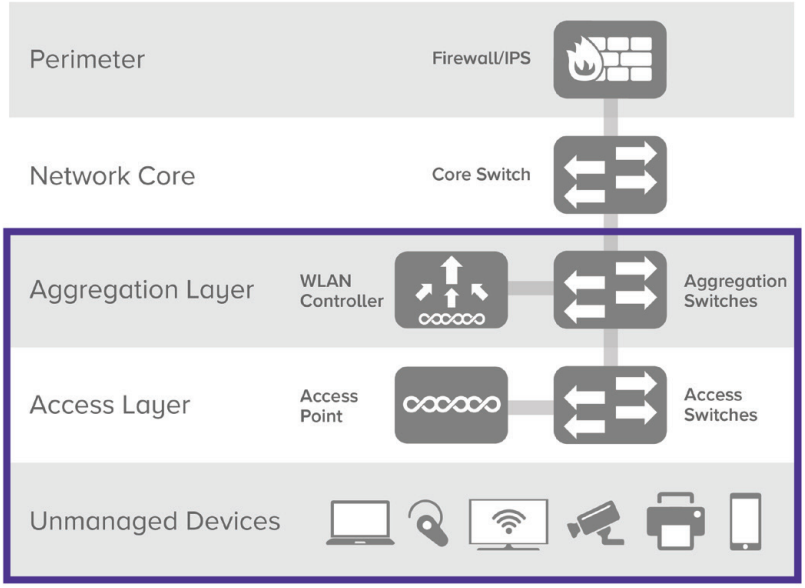

Flexibility and Integration

While visibility into the wireless access layer will be essential, it will often be valuable to correlate and track behaviors deeper into the network as well. Thus it will be beneficial for a solution to collect information higher in the network. Again this can be done in a variety of ways either through dedicated sensors or via integration with existing packet capture solutions or perimeter devices. Network architectures will naturally vary, so the key will be to provide flexible options that can add additional coverage while keeping costs and complexity in check.

Easily extend visibility and add context with additional data sources.

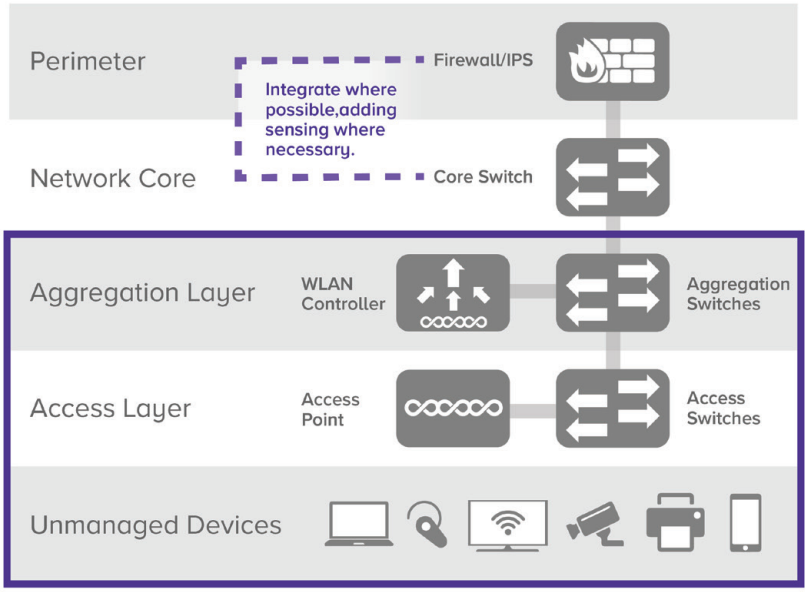

CONCLUSION

Just because a device is “unmanaged” in the traditional sense does not mean that it is unmanageable. With the appropriate architecture, security teams can still monitor the behavior of any device, enforce policies on what it is allowed to do, and block malicious actions. But this goal comes with new requirements. Since we can’t control the devices themselves, we will need to improve our control over the environment they operate in. This means directly monitoring the wireless environment, tracking all devices, controlling what connections are allowed, and inspecting traffic for threats.

While these functional requirements have long been taken for granted at the network perimeter, they have been absent in the deeper layers of the network. As more and more devices include a variety of network options by default, and the traditional perimeter becomes more fluid, it is imperative that organizations adopt new strategies that extend visibility to all devices and all corners of an organization. While this may sound daunting, the path to control is not as steep as it may initially sound. In most cases, we can simply add new layers of intelligence and threat detection to the wireless and wired infrastructure that already exists. In this way we can build an architecture that is prepared for the realities of the modern device and security landscape.

Sources

[1] IDC, Rapid Maturity of IoT Projects Highlights Risks, July 2017

[2] Tech Republic, Report: IoT attacks exploded by 280% in the first half of 201, August 2017

armis

1.888.452.4011

armis.com