Understand The Basics of Cybersecurity

...and be comfortable in a security conversation.

By Trevor Pott

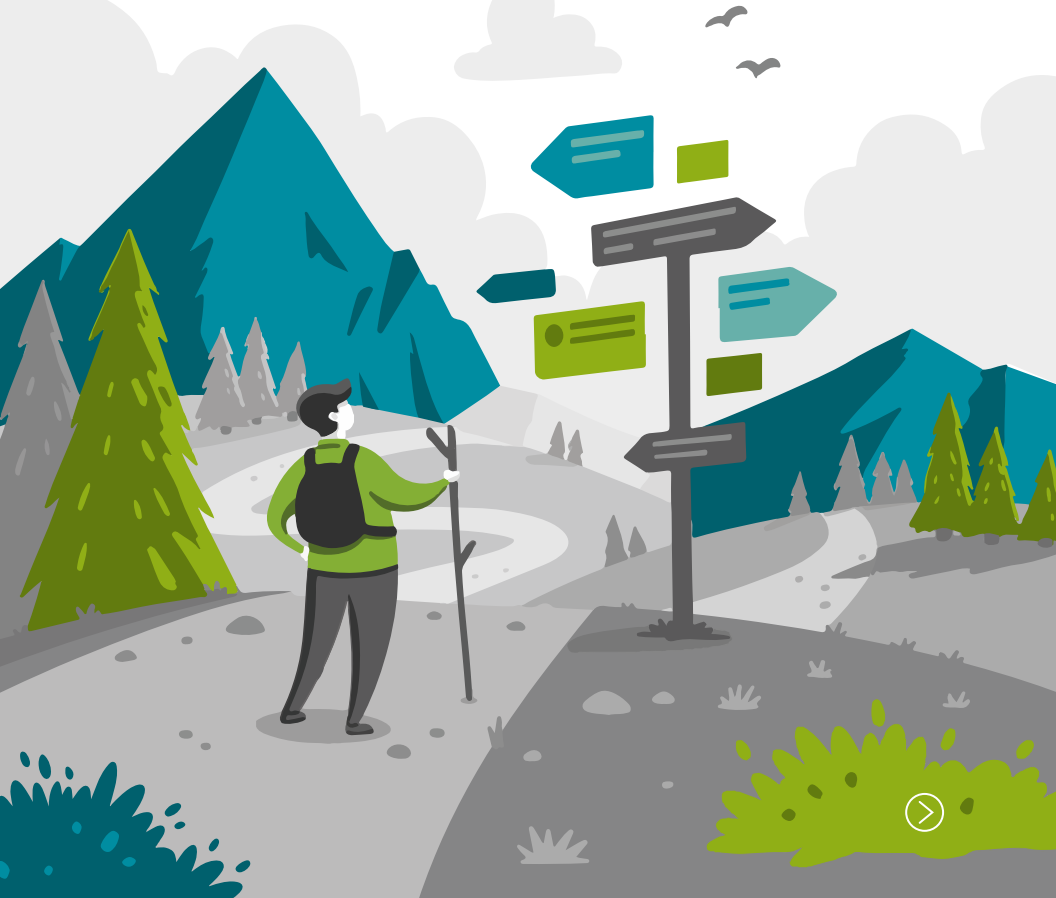

CONTENTS

01 Intro to infosec

5 Infosec, referencing birds
7 No computers required
8 Ignorance
10 Social Engineering
11 Negligence
13 Winter is coming

02 Basic IT infosec concepts

17 Infosec is about much more than firewalls and anti-malware
19 Doing infosec like it's 2005
23 WAF and Application Security
25 Encryption and DLP

03 Basic IT infosec defenses

29 Monitoring
32 DLP and CASB
34 SIEM and ATP
36 Access Control, VPN, and Remote Access
39 Browser defenses
40 MDM
41 Authentication
42 Automation

04 Anatomy of an attack

46 Defining the objective
50 Connectivity
54 Paying
57 Reconnaissance
60 Getting to know the target
63 Thinking like a hacker

05 Advanced concepts

67 Regulatory Compliance
68 Network Segmentation and Microsegmentation
70 The Internet of 5G Things
72 Putting it all together

06 Juniper connected security

76 Security needs layers

INTRODUCTION

Information security, often shortened to “infosec”, concerns the hypotheses, tools, techniques, technologies, and practices surrounding the access to – and protection of – information. Information security is not limited to specific products or services. Infosec should instead be thought of as a way of thinking: a set of goals regarding information, its access and its usage, that technology can help us achieve.

Many of you who read this will no doubt be looking for a color-by-numbers guide to information security. I’m sorry to inform you that such a thing does not – and cannot – exist.

There is no series of tick boxes that one can follow to secure a network. There are no magic buzzwords to sell infosec more easily, nor are there universal use cases to be cheaply and easily exploited by marketing.

Information security is constant evolution. Information exists in our minds, on paper, in our computers, and even exists to be gleaned from our actions. The simple patterns of our daily lives are useful, exploitable information, which is why companies like Facebook exist, and make the billions upon billions of dollars every year that they do.

The purpose of this document is not to provide you with a magic talisman that will somehow make you an infosec rock star. Instead, the aim is to teach you the fundamentals, so that every time you encounter a new infosec technology, you are able to understand why it is useful, what it defends against, and how it helps counter the aspects of human nature which ultimately make our networks vulnerable.

Trevor Pott

Technical Security Lead, Juniper Networks

01

INTRO TO INFOSEC

INFOSEC, REFERENCING BIRDS

Not only does infosec concern more than just the information on computers, the adoption and refinement of both practices and tools related to information security aren't even limited to our species. There are a number of examples within the animal kingdom of species that practice infosec, each with varying degrees of complexity and success.

Consider the humble nuthatch. Nuthatches are a genus of birds – Sitta – weighing between 10 and 47g, depending on the exact species. Nuthatches can be found in many climates, and do not generally migrate. In cold climates, such as Canada, nuthatches must cache food for the winter.

Nuthatches will hide seeds, berries, and other food items underneath the bark of trees. Over the course of a summer, a single nuthatch will cache thousands of individual food items, each in their own individual hiding place.

Nuthatches flock with other small bird species, such as chickadees. All birds in the flock will gladly steal food from one another, especially in the depths of winter, as will squirrels, and birds from rival flocks.

To survive, nuthatches must be circumspect. When caching food they will often fly away from the food source in a different direction from where they ultimately intend to cache that food item, frequently flying circuitous paths in order to confuse observers. Nuthatches also try to hide their food in locations where only a nuthatch, with their specialized beak, can get at the food.

Other bird species work to counteract these defenses. Some track nuthatches by calling out when one is seen, helping the entire flock of that species keep track of where the nuthatches are going after picking up food from a food source. Some birds use sticks, pins, or other tools to extract food items hidden by nuthatches in places their beaks can't access. In response, nuthatches attempt to be more clever about where and how they hide food: their survival depends on it.

Nuthatches have data they need to keep secret. This data is the location of their winter food caches, and the means of extracting food from these locations. By the time winter rolls around, every other animal is attempting to find that data, so they can steal the nuthatch's food. If the nuthatch fails at infosec, they die…or they become hackers themselves, stealing data – and food – from others.

This is information security, boiled down to its purest essence.

No computers required.

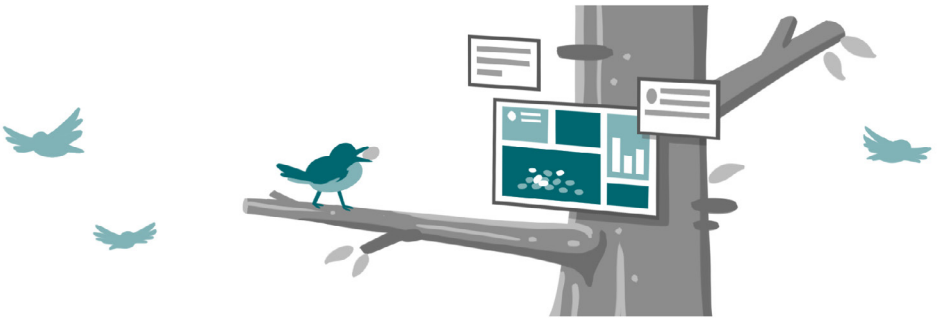

NO COMPUTERS REQUIRED

In today's world, most information is stored in computers. As a result, infosec is often focused on IT. It is, however, important to recognize that information security begins and ends with people. People have information to secure; other people are going to attempt to uncover, alter, or delete that information; and people are the greatest source of information security vulnerability.

This last point – that people are our greatest infosec weakness – is absolutely critical to understand. Compromise events that use previously unknown vulnerabilities to break into networks in a novel fashion are almost unheard of. In the real world, information security compromises happen almost exclusively because a human being responsible for defending data failed to do so.

Technology can be used to prevent information security compromise events by directly blocking various attack types, offering mindfulness reminders to authorized human operators, or attempting to detect either malicious, abnormal, or negligent behavior. Broadly speaking, there are three categories of human compromises: ignorance, social engineering, and negligence. Each of these is distinct, though they are often mistaken for one another.

IGNORANCE

Ignorance is, quite simply, a lack of knowledge. Many societies around the world – but especially modern Western societies – stigmatize ignorance. Popular depictions of knowledgeable persons, especially scientists and technologists, frequently show these individuals possessing an impossible amount of knowledge. Think Scotty from Star Trek.

We are raised to believe being Scotty is just how "nerds" are. Being a "nerd", and not having Scotty-like omniscience is not only abnormal, but we are frequently told by our culture that any display of ignorance is immoral as well.

Unfortunately, not only is being Scotty impossible, but for some reason we tell our children that not being Scotty is bad, wrong, and makes one a failure. We raise them to believe that failure is a character flaw, and then we go tell them to defend our networks against unknown threats, using rapidly evolving technologies, under unrealistic budget constraints, and with impossible deadlines.

Our collective antipathy towards ignorance can be seen in many aspects of our culture, including our humor. Consider, as an example, the website https://lmgtfy.com/, more popularly known as "Let Me Google That For You".

When an individual demonstrates ignorance by asking a question, someone else will enter the question into LMGTFY. This creates a link with a sarcastic animation that shows the person who asked the question how to search for the answer to their question using a search engine. The entire purpose of LMGTFY is to mock people for being ignorant, and it is only one amongst many readily available examples.

To be perfectly clear, so that there is no possibility of confusion: there is absolutely nothing wrong with ignorance. We are, all of us, each and every one of us, ignorant. Even the brightest among us are ignorant of far more than they will ever know. This is part of the human condition. Our brains only hold so much, and we have only minimal control over what gets crammed into our noggins for long-term storage.

There is a quote that I, your author, am rather fond of: "spotting an expert is easy: they're the person who, when asked to explain their field to you, constantly qualifies every statement and explains the limits of the models they are using". In other words, an expert is not someone whose knowledge on a given topic is complete. An expert is someone who knows enough to understand just exactly how much they don't know.

Where ignorance becomes a security risk is when we believe that we know something, when in reality we do not. Hubris – excessive pride or self-confidence – is what makes ignorance problematic.

Ask an infosec expert to defend your network, and they will immediately bombard you with questions. The overwhelming number of these questions will not be about technology. Beyond simple needs assessment, their questions will include topics such as your risk appetite, the types of data you are defending, the regulatory environments in which you operate, and more.

Experts know what they don't know, and in order to do their jobs they must overcome our societal taboos about ignorance and actually ask questions. Every question left unasked means that an assumption must be used in place of an answer. Every assumption made is a potential vulnerability to be exploited.

If you're a nuthatch, you don't just assume nobody is watching when you fly off to store your food; you look around and check, and then throw a bit of a flourish into your flight pattern, just in case. Hubris is the quickest path to ensuring one's own downfall, whether you are trying to defend a computer network, or just trying to survive the winter.

SOCIAL ENGINEERING

Social engineering is a polite way of saying “pulling a con”. To socially engineer someone is to convince them to do something you want them to do, and that something is almost always not in the mark’s best interests.

In the context of IT, popular scams include the simple calls from Sharon, your local Google specialist or Windows tech support, or spear phishing emails on the sophisticated end. A spear phishing email might be carefully crafted to appear to come from your bank, your boss, or your CEO. They typically attempt to get you to reveal your password, often by convincing you to click a link, and then log in to a website that – if done right – looks absolutely identical to the real one.

Social engineering isn’t only a tool of the attacker. Our friendly neighborhood nuthatch uses it too. Some percentage of the food cached by a nuthatch is empty seeds or rotten berries. This food is often cached with minimal effort put into hiding the location. Alternately, it might be cached with regular amounts of obfuscation, but used as a distraction when it is time to eat.

The goal is to fake out observing birds by convincing them that the delicious goodies are somewhere they are not. While they’re occupied with the distraction, the nuthatch goes about its business unmolested. In IT, a similar concept is that of the “honeypot”: a deliberately vulnerable system festooned with monitoring to entice the bad guys into attacking it, and in doing so revealing themselves.

All of us have fallen prey to a scam at some point or another. Even the most paranoid among us cannot overcome our very human need for interpersonal connection, and this makes us vulnerable. This vulnerability is ruthlessly exploited by hackers, and exploiting this vulnerability is a skill that information security professionals on all sides of the battle actively spend time levelling up.

NEGLIGENCE

Negligence is a failure to do something that you should have done. Ignorance and social engineering are often blamed on negligence, but there are some important differences. To start with, all of us are ignorant, and all of us are vulnerable to social engineering. To contrast, negligence is a choice.

Each and every one of us has been scammed. We will all be scammed in the future. There will always be far more in this world that we do not know, when compared to what we know, even if we narrow the scope of the discussion to a field in which we are supposedly experts. This is simply part of what it is to be human, and all the social taboos in the world can't change that.

These topics may be uncomfortable for many, but they are important. Because of our social taboos, we often accuse those who are victims of social engineering of being negligent, or conflate ignorance with negligence. They are not the same thing.

Negligence requires someone to deliberately choose to put aside some aspect of their duty, whether that duty is educating themselves, or performing a task. Negligence is also distinct from laziness, although laziness can lead to negligence.

Laziness can manifest itself through useful pathways. A lazy systems administrator, for example, might automate every mundane task they encounter, so as to not have to repeat them. This expression of laziness is useful. Laziness can also manifest as negligence: choosing to forgo education, business practices, or necessary administrative tasks because they are annoying, burdensome, or so forth.

In addition to taboos around ignorance, social engineering and laziness which encourage us to classify these as purposeful moral failings – and thus negligence – there are also social pressures to blame "expendable redshirt" IT practitioners for all compromise events. Sysadmins may miss critical steps not because they are negligent, but because they are overburdened.

When sysadmins are overburdened, negligence may be the result of sysadmins not speaking up and explaining that they are swamped with too much work. All too often, however, the real negligence lies with their managers, who simply refuse to provide the necessary tools, manpower, or time off. Sysadmin burnout is a very real problem, and the quality of work degrades sharply before sysadmins flame out and crash.

WINTER IS COMING

What is important to remember is that while the metagame between attacker and defender is constantly evolving, the basics of the game haven't changed for quite literally billions of years. Information security is information security, whether you're a network administrator, or a nuthatch.

Understanding information security begins with understanding people. What motivates us? How can we be convinced to do something? How can we be tricked into making a mistake? Which mistakes are defenders most likely to make, under what circumstances, and when?

The products and services sold by IT vendors are just tools: a new place to cache your seeds, or a new tool to pry that same cache open. Understanding the tools of the day will help you fight today's battles, but understanding the underlying principles of information security is what's needed to survive the war.

02

BASIC IT INFOSEC CONCEPTS

With the universal basics of infosec out of the way, it's time to focus on more IT-specific concepts. After all, the overwhelming majority of human information is stored in a computer somewhere. In the context of IT, information is more than just "data". Or, perhaps more accurately, the information that systems administrators and information security professionals must concern themselves with is more varied than what most of us think about when someone says "data."

To most of us, the word "data" evokes thoughts of files, or databases. Pictures, word documents, customer records and so forth are all common types of data that most of us interact with on a daily basis. Information about information – otherwise known as metadata – is also important, as is monitoring data.

Consider a picture taken from a smartphone, and then uploaded to a social media website. Pictures often have metadata associated with them: information like GPS coordinates, the model of the phone or camera that took the image, and more.

If this metadata is not stripped out when it is posted to social media, then a malicious actor can examine the image, and determine a target's location, and even what type of phone to look for in order to help identify that individual. There are numerous documented cases of thieves using this data to burglarize someone's house.

Consider someone taking a food selfie, and attaching it to a social media message saying "look at what I am eating". This gives a thief both a location and a time; simple maths will tell them whether or not that individual can get back home before the thief can steal from them.

Metadata and monitoring data can reveal carefully guarded corporate secrets as well. An electrician posting a picture about the really nice-looking cabling job they just completed on the new as-yet-unannounced datacenter could easily give away the location of this datacenter via the metadata in their pictures. Similarly, a Word document tends to keep a history of everyone who has edited it, and this can have legal consequences.

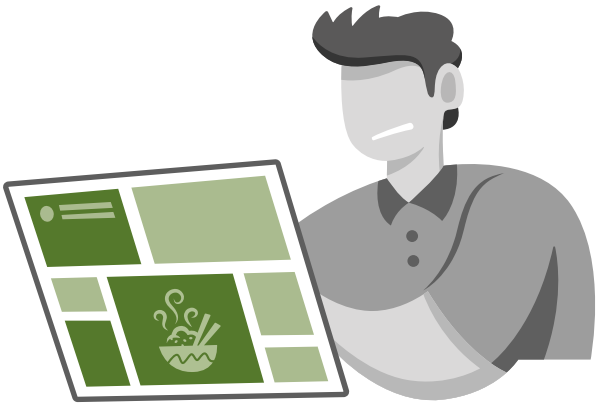

INFOSEC IS ABOUT MUCH MORE THAN FIREWALLS AND ANTI-MALWARE

If the layperson thinks about information security, the first thing to mind is often a firewall, or an anti-malware application. Home networks are often defended only by the firewall built into the router that their Internet Service Provider (ISP) provides them, as well as the anti-malware software that is built into their computer's operating system. The rare consumer has a third-party anti-malware application.

A classic firewall, such as that found on a home router, concerns itself with data passing between the internet and one's home network. Most consumer firewalls are extremely primitive, providing security predominantly through a default ruleset that says "nothing can send data to this network unless a device on this network has asked for that data".

Consumer firewalls are primitive in part because they are configured to allow all traffic out. The other part of why they are considered primitive is that they perform virtually no inspection on inbound traffic, beyond ensuring that the IP address and port is configured to allow traffic. Consumer firewalls are designed not to require customer interaction, and this limits how secure they can be.

The majority of consumer anti-malware is also fairly primitive. Anti-malware concerns itself with files, looking for signs that a file is not what it appears to be. If an individual opens or saves a file, the anti-malware will scan the file. If the file matches a known malware signature, or has characteristics that are very similar to known malware, access to that file will be denied. Many anti-malware applications also run periodic scans in order to detect malware that is stored on a computer, but not currently being accessed.

Classic firewalls and consumer anti-malware are good models to understand the two primary areas of concern for information security: data in flight, and data at rest. Data in flight refers to information that is transiting a network. Data at rest refers to information that is stored in a permanent fashion, such as on a hard disk.

In enterprises, the majority of information security products and services focus their efforts on data in flight. Through any of a number of means, the goal is to catch whatever badness is happening before it reaches the user, or leaves the network.

In an enterprise, infosec products focused on data at rest tend to focus more on "is this data backed up" and "can this data only be accessed by the individuals or applications that should be accessing it". There are a few products that focus on anti-malware, but they are a small fraction of the products aimed at data at rest.

Consumer networks don't tend to have anti-malware capabilities built into their data in flight defenses. Enterprises do (or at least, they really should). In an enterprise, anti-malware that relies on catching malware by examining data at rest is considered a last resort. If the anti-malware program on your PC catches an infected file, then something has gone very wrong with the umpteen layers of information security products and services that should have stopped the unspecified badness from ever reaching your PC in the first place.

Enterprises use infosec products beyond simple firewalls and anti-malware for a number of reasons. The first is that both classic firewalls and data-at-rest-focused anti-malware are not particularly effective.

A classic firewall can prevent someone from, for example, remote controlling your PC using a direct connection to an Internet Protocol version 4 (IPv4) address. It cannot filter out phishing emails that contain a link which, if clicked, will either download a remote control application designed to work from behind a firewall, or which would reconfigure your firewall to allow someone to remotely access your PC.

Similarly, anti-malware applications are quite good at defending against malicious files of a known type, but they're less effective against malicious files that the anti-malware vendor hasn't seen before, and virtually useless against the myriad interesting things that can crawl onto your PC through an internet browser.

DOING INFOSEC LIKE IT'S 2005

Sometime in the early 90s – exactly when depends on whom you talk to – the first layer 7 (application layer) firewalls were created. Classic firewalls operate at layer 2 and 3, and simply block network protocols or IP ports. If you want to allow someone to access a service located behind a classic firewall, you simply open the port, and allow inbound traffic directed at that port to reach an application located at a specific port and address. Home users will likely be familiar with this as "port forwarding".

To contrast, application layer firewalls are (to varying extents) application-aware. Application-layer firewalls defend networks by acting as a proxy between an application and the outside world. In this case, when a request is made for access to an application behind the firewall, the request is not transparently passed through.

An application-layer firewall will, at the very least, inspect some part of the traffic involved in the request, with an eye to ensuring that the request is legitimate, and not malicious. Modern application-layer firewalls will act as a proxy, in large part because HTTPS streams cannot be inspected without a proxy to perform decryption, and today, web traffic is almost all encrypted.

A proxy lives between the application and the requestor. When a proxy is involved, the requestor's request ends at the proxy. The proxy then makes its own request back to the application. In this way requests can be more carefully screened for malicious activity. Proxies can also deny certain types of requests whilst allowing others.

The most application-aware of the application-layer firewalls know all the appropriate commands that an application can receive, and thus can ensure that no commands are sent to the application which would disrupt it. This level of application awareness typically includes multiple layers of badness checking, the details of which vary from application to application.

Somewhere in between the early application-layer firewalls and the more proxy-based firewalls that became common in 2005, firewalls ceased simply being firewalls, and became Next-Generation Firewalls (NGFW). NGFWs are more integrated into organizational infrastructure, and bring advanced capabilities, such as the ability to apply policy based on username or group, rather than just by IP.

Exactly what an NGFW is, or who built the first one depends entirely upon whom you ask. In general, a NGFW is one that not only prevents bad guys from getting into a network from the outside using basic attacks, it also offers some form of data-in-flight scanning capabilities.

A good example of a modern NGFW capability that is common (but where vendors will disagree with nomenclature) is the ability to scan emails that pass through it, looking for phishing attempts. For some vendors, this functionality was reserved for discrete Threat Management (TM) and/or Unified Threat Management (UTM) products. Other vendors built this capability directly into their NGFW products.

The majority of consumer firewalls do not have anti-phishing capability, nor do traditional anti-malware applications. This is problematic, because people are incredibly vulnerable to social engineering, which is what phishing emails rely upon.

Both Next-Generation Anti-Malware (NGAM) – called Next-Generation Anti-Virus (NGAV) by some vendors – and NGFWs can usually be distinguished from their predecessors by the ability to help defend against social engineering attacks. Most NGFWs also offer legitimate users ways to connect into the network from the outside, while NGAV products tend to include a host-based NGFW.

To note: the existence of NGAV does not obviate the need for a NGFW. They operate at different levels. NGFW tries to stop the badness from getting to your PC in the first place, which is the best, and safest option. NGAV tries to prevent the badness from compromising your PC once it has reached there, and as such is a last line of defense. In addition, NGFWs can protect thousands of users at once, and can defend endpoints incapable of using NGAV (such as printers) whereas NGAV can only defend the host upon which it is installed.

Commercially viable NGFW and NGAV products were common in 2005, and table stakes for enterprise infosec by 2010. Infosec has evolved considerably since then, in large part because of the widespread adoption of public cloud computing.

WAF AND APPLICATION SECURITY

The evolution of the application-layer firewall is the Web Application Firewall (WAF). A WAF is an application-layer firewall with extra application-awareness blue crystals for that deep-down clean that all data in flight really needs. As one might expect from the name WAFs are focused on HTTP and HTTPS-delivered applications, typically ones that are accessed from across the internet, using the World Wide Web (WWW).

To be extra confusing, however, there's nothing that says a WAF must use the WWW, or any other part of the internet. There are plenty of use cases for plopping a WAF in front of a web server that will only ever be accessed by other applications on the same network.

WAFs protect against vulnerabilities specific to a given application, or class of application. The most famous example is protection against SQL injection attacks (obligatory XKCD). An SQL injection attack happens when someone fills a field intended for another type of data (typically a text box on a website) with SQL commands.

If the application isn't specifically designed to strip SQL commands out of fields where they aren't supposed to appear (an approach called input parsing), then those SQL commands can execute. In this way, an attacker might cause a database to delete itself simply by entering some SQL commands into the name field of a web form. A WAF would filter out SQL commands, along with many other forms of attacks, with the attack types varying based on the individual web application.

WAFs are often used internally to protect against insider threats, as well as to protect one layer of an application from being affected by the compromise of another. If a load balancer were compromised, for example, having a WAF between it and the web server itself could prevent very bad things from happening.

Input parsing is a type of application security. Application security does not always rely on firewalls, anti-malware, or any other external product or service to provide security. Traditional application security builds layer upon layer of protection directly into the application itself.

Like anti-malware applications on a PC, application security built directly into the application is typically viewed as the absolute last line of defense for server-based applications. Especially in enterprises, if attacks make it far enough that the application's own security features are providing defense, something has gone horribly wrong with all the other layers of defense that should exist between the requestor and the application.

ENCRYPTION AND DLP

Encryption and Data Loss Prevention (DLP) occupy the other side of information security. Instead of focusing on preventing attackers from breaking in from the outside, they are concerned with preventing data that's already inside the network from getting out.

Encryption technologies ensure that only individuals or applications with the correct key can access data. Encryption technologies exist to protect data in flight, as well as data at rest. Encryption is used to defend data at rest against scenarios such as the theft of a device, or someone compromising a virtual or container host.

In these scenarios, if the data is unencrypted, the attacker would have access to all the data on the target device, Virtual Machine (VM) or container. Encryption for data at rest is increasingly mandatory in organizations of all sizes. A lot of damage can be done to an organization if, for example, a sales manager's laptop is stolen, which happens to contain a complete list of the organization's most valuable customers, and all data related to those customers.

Encryption on data in flight is used to ensure that data is not "snooped" upon by third parties. This is especially important when accessing internet-delivered services, or when accessing anything over a wireless network.

Neither the internet nor wireless networks can ever be made 100% secure. As a result, encrypting all data in flight (such that only the requesting party and the application or service they are trying to access can obtain the information being exchanged,) is absolutely vital to modern IT.

DLP technologies are intricately linked with both NGFW and NGAV. Each individual implementation of DLP is different, but the fundamental approach behind DLP is to scan data that is attempting to leave the network, and to freak out and stop it if it is not supposed to be leaving the network.

DLP technologies that scan emails and files are fairly mature. NGAV applications are increasingly building file-based DLP into their offerings, preventing files with certain kinds of content from being copied to a flash drive, or placed in a cloud storage folder (such as Dropbox).

DLP is also being incorporated into many NGFWs as a standard feature, but as an emerging technology is sometimes incorporated into other product categories instead. Here, DLP is usually focused on scanning email traffic passing through the firewall to see if something is being transmitted that shouldn't be. NGFW-based DLP may also contain cloud-storage analysis capabilities, and might prevent data from being stored on a cloud storage service by analysing the files as they are being transmitted to the cloud storage service.

DLP is currently the bleeding edge of infosec. Most technologies in this area are nascent, and there are a lot more ways to sneak data (exflitrate) out of a network than simply copying files to a USB key, using a cloud service, or emailing them.

Instant messengers, IRC, Slack, and social media are all exfiltration vectors that are poorly covered by modern DLP. For many data types, there is also the option of simply pointing one's cell phone at the screen, taking a picture, and then walking out of the building.

The infosec war between attacker and defender is eternal. The technologies change, but the basic concepts don't. At the end of the day, information security professionals try to keep the bad guys from knowing information they shouldn't, while trying to prevent the bad guys from sneaking in information that's harmful.

03

BASIC IT INFOSEC DEFENSES

To understand how to defend against attack, one must understand the tools available. Firewalls, anti-malware, encryption and DLP have all been explored, however, these technology categories should be considered more akin to building blocks than actual products today.

Information security relies on a concept known as defense in depth, which is effectively the industry term for layering defenses one atop another, in the hopes that if enough different kinds of defenses are in place, the majority of the badness will be prevented. Defense in depth is not a new concept, and infosec vendors are regularly incorporating new functionality into existing products, or increasing the interoperability of their products such that chains of independent products can provide defense in depth.

In the long run, interoperability is the better strategy. No single vendor can defend a network against today's threats. No vendor has enough manpower, or enough research and development capacity to chase innovative new approaches to counter emerging threats. The only realistic way to defend a modern network is to have multiple products from multiple vendors working tightly together.

Understanding what the various commonly overlapping infosec technology categories do can help make a determination as to whether feature overlap is an unnecessary waste, or a useful reinforcement of defenses.

MONITORING

Monitoring enables IT at scale. Without monitoring, IT practitioners have no way of knowing what is occurring, much less responding to problems. Monitoring data is also a double-edged sword.

Monitoring data is useful to pinpoint compromises. It is useful to attackers as well. Without getting too deep into the technical weeds, the real purpose of modern monitoring applications is to spot patterns. Those patterns can tell IT practitioners when something has gone wrong, but they can also serve as neon signs pointing attackers to where the interesting stuff lives.

The majority of today's monitoring products spot patterns using one of two approaches: deviation and correlation. Deviation relies on observing workloads and data flows between systems that are behaving as expected in order to determine what "normal" looks like. This process is known as baselining.

Baselining can be triggered manually by IT practitioners, or in more advanced products, can be done in an automated and ongoing fashion. Establishing a baseline is reasonably straightforward: observe various data points, figure out how much the various metrics wiggle around, and in what pattern.

As a general rule, when something deviates from baseline, one of three things is happening: there has been an application or hardware failure, workloads are under stress and need to be scaled, or there is a compromise event taking place. Anything that deviates significantly from the established baseline is clearly doing something abnormal, and should be looked at. Exactly how deviation from baseline is determined, as well as what the rate of false positive and negatives are is the special sauce that differentiates one baseline-based monitoring product from the next.

Baseline-based monitoring typically uses real-time monitoring data, and as such is often both resource-intensive, and difficult to scale. This can be mitigated by only sampling data periodically – a few seconds out of every minute, for example – but the lower the polling frequency, the less accurate the results.

Correlation is a far more common approach to monitoring. Correlation typically relies less on real-time monitoring data, and more on the event logs generated by individual applications, or infrastructure components. This isn't a hard and fast rule; many correlation-based systems do poll real-time data, however, the monitoring data history kept by these products rarely goes back very far.

Correlation-based monitoring looks for multiple events occurring either at the same time, or in rapid succession. Often, a correlation-based monitoring system looks for a series of events that are known to occur together that indicate a problem.

As with baseline-based monitoring, correlation monitoring is often used to detect equipment or application failure, as well as compromise events.


1
2
3

DLP AND CASB

Monitoring, in one form or another, is at the heart of many advanced information security products. The incorporation of monitoring data into advanced infosec products is usually not quite so straightforward as simply baselining or correlation, and regularly draws on multiple techniques to achieve the desired goal.

DLP, which was discussed in the previous chapter, is an example of a technology category in which multiple different approaches to solving the problem exist. Some of these solutions rely on monitoring, while others do not.

When monitoring is part of DLP it is often baseline-based, and focused on deliberate insider threats. The ideal monitoring-based DLP system would examine a user's data access of multiple data types across the entire organization's IT infrastructure, looking for abnormalities.

If, for example, a sales manager who normally only accesses information on a handful of accounts in a given day suddenly pulls information on all accounts in their region, something untoward might be afoot. This approach to DLP is bleeding edge technology, and requires multiple products from multiple vendors to work together.

Total DLP coverage of all organization data is often impossible. Enterprises, for example, can have data in thousands of different organizations, both on premises and in the public cloud. Even if an organization could monitor all data access, there is a question of what to do about it.

Abnormal access patterns could be an employee trying to steal data, a compromise event occurring because of an external attacker, or just someone trying to do their job. Increasingly, machine learning is being used to fine tune these parameters, and this is expected to be a product category that sees significant development through the early part of the 2020s.

The other primary approach to DLP is scanning content, and placing restrictions on it based on context. In order to accomplish this, DLP is often incorporated into various security products, including Cloud Access Security Brokers (CASBs), and/or Advanced Threat Protection (ATP) products. This approach to DLP is more mature, having been used in a variety of products at enterprise and even service provider scale for years.

Content scanning DLP systems have a number of approaches. For some file types, such as spreadsheets, the contents of the file itself are scanned. If a spreadsheet with credit card information is encountered, this can generate an alert, or be blocked. CASBs are often used to enforce policies about what data can be uploaded to which public cloud services, or be combined with other policy enforcement products to restrict which cloud services can be accessed at all.

SIEM AND ATP

Security Information and Event Management (SIEM) products are the hub of information security, and are increasingly interwoven with ATP products. In the early 2000s, SIEM was simply a marketing term for monitoring products, and ATP has traditionally been a marketing term for network defenses which scan data flows for malicious email attachments, attempts to access restricted websites, or calls by malware to command and control servers.

Over the years, the two categories of products have evolved, with significant feature overlap between them. Both SIEMs and ATPs serve as recipients for data from a number of products. The focus of which types of data are ingested differs between the two categories of products; however, the gap is closing as the years go by.

The products sending data to the SIEM or ATP can be a collector or agent that is related to the SIEM or ATP in question. They can also be a third-party product such as a NGFW, CASB, anti-malware, and so forth. Most SIEMs and ATPs do not rely on collecting data only from products supplied by the same vendor, and the most useful ones have simple and quick integration with leading products from multiple vendors.

The most successful SIEMs and ATPs are strongly focused on correlation. As an example scenario, let's consider an ATP which is integrated with both a NGFW at the network edge, the NGAV on client systems, and a CASB.

The NGFW sends email data streams to the ATP, which detects a series of malware-laden emails targeted at a given user. The NGFW is instructed to prevent the emails from arriving, and does so. Shortly thereafter, the NGAV on that user's endpoint detects strange behavior. In rapid succession, the CASB detects that the endpoint in question is rapidly attempting to connect to cloud storage sites and upload documents.

While individually these items might raise an alert, the severity of the alerts for any of these individual events is likely to be low. Combined, however, they tell a more concerning story. The ATP would correlate these events and determine that a targeted attack against a given user was likely underway.

The initial round of malware was blocked, but clearly something got through. That something made it past the NGAV, which detected something unusual, but either couldn't do anything about it, or the unusual behavior didn't reach the threshold required for the NGAV to prevent the activity. The unspecified badness that managed to make it onto the now compromised endpoint was detected by the CASB as attempting to exfiltrate data using cloud storage. This series of events would cause an ATP product to declare this a max threat, and call for a human.

ATP products incorporate badness scanners that come in flavors. Other products, such as NGFWs can pump data through an ATP, and the ATP will scan it for badness. Many ATPs, however, can also ingest event and even monitoring data from other products, and this is what creates the overlap with SIEMs.

SIEMs are predominantly focused on collecting event and monitoring data. It is what they were originally designed to do. Because SIEMs are already the hub of information for monitoring and event data, it wasn't a big stretch for many SIEM vendors to start grafting badness scanning capabilities into their products.

Because ATPs are more focused on badness scanning – whether that be scanning attachments, DNS queries, HTTP streams, or what-have-you – they tend to assist with DLP predominantly by integrating with NGFWs and/or CASBs, and using a content scanning method. When SIEMs participate in DLP, it is usually by either having monitoring-based DLP built-in, or integrating with an application that does.

ACCESS CONTROL, VPN, AND REMOTE ACCESS

Access control is an IT industry term for the somewhat nebulous and vaguely defined concept of limiting access to IT resources. The concept is somewhat nebulous because exactly what constitutes "access" changes depending on the person, organization, or law attempting to define it.

Thanks to infosec technologies like encryption, it is possible to have physical access to the storage device that contains a given dataset, but not have access to the data. Similarly, cloud computing can allow organizations to have access to products, services and data from anywhere in the world without ever having access to the underlying hardware.

To many IT practitioners, access control is synonymous with Access Control Lists (ACLs) which are a way of codifying which users and/or groups have access to which resources. Other administrators will see access control as limiting physical access, while still others will look at the problem through viewpoints ranging from automation, to service brokers, to data sovereignty.

Give the best infosec people in the world physical access to a server with an encrypted hard drive, and under the right circumstances, they can find the encryption key. Thanks to side-channel attacks, under even more stringent circumstances an individual compromising a virtual machine can access data from other virtual machines to which they don't have what would traditionally be called "access". The technical details of attacks change every year, but differing viewpoints about what constitutes access are persistent.

Virtually every piece of IT infrastructure, every operating system, and every application has some form of access control. Among the two most important methods of access control are Virtual Private Networks (VPNs) and Remote Access.

VPNs are an encrypted network tunnel between two computer systems. VPNs are used to allow individuals to connect securely to their organization's network, by individuals in repressive countries looking to escape censorship or work around region locks for online content, as well as to establish secure links between sites, especially over the internet.

VPNs have two primary functions. The first is that VPNs can allow access to the private address space used by an organization. Private addresses are not routed across the internet, and thus cannot be accessed without a network tunnel.

The encryption used by VPNs offers protection against snooping by malicious actors. The internet is not a safe place, and it is not unheard of for one or more of the routers, switches, servers, or other systems through which our data packets flow to be compromised. Compromised network equipment can be used to spy on unencrypted data streams, which is where VPNs come in.

While Remote Access may sound as vague a concept as access control, it's not. In fact, it's an important enough concept to earn it capitalization. Remote Access is the term used for any of a number of technologies that allow individuals to access organizational resources without using VPNs.

While VPNs have performed admirably for decades, they have a number of drawbacks that make their use burdensome and unpopular. Speed issues, finicky VPN clients, and/or difficulties implementing mitigating technologies such as split routing and split DNS mean that evidence has been mounting for some time that people prefer access to organizational resources without VPNs wherever possible.

Both VPN and Remote Access technologies are the primary form of access control used by organizations to defend their networks. VPNs and Remote Access allow employees to access organizational resources that are protected by the organization's primary infosec defenses.

Traditionally, security behind the perimeter is less restrictive than it is regarding attempts to access resources from outside the perimeter. This is largely because IT workloads and infrastructure have to be able to interoperate and interact with one another.

The perimeter model is used because it would require tremendous manpower, money, automation, and possibly all three to adequately defend applications designed to operate in a perimeter model, while having them exposed directly to the internet. This problem is being addressed by emerging infosec technologies such as CASBs, ATPs, and microsegmentation. Augmented by the adoption of next generation Remote Access products, the traditional network defense perimeter model is evolving, and the role of VPNs along with it.

BROWSER DEFENSES

Other than email, web browsers are probably the most likely means by which an individual user's endpoint can become compromised by an external attacker. Today's web browsers are more secure than those of yesteryear, but still quite vulnerable. Browsers also allow end users to download files from the internet and then execute them, the third most common means of compromise.

Popular web browsers, such as Chrome and Firefox have the ability to install browser extensions. Popular browser extensions such as Adblock, Ghostery, and Privacy Badger offer additional protection against various forms of internet badness, such as malvertising. These types of browser extensions aim to prevent the web browser from attempting to request a connection to suspicious internet resources.

Other browser extensions, such as those provided by NGAV vendors, as well as CASBs, NGFWs, ATPs, and carefully curated DNS servers all attempt to protect end users by preventing a request to access a compromised resource from completing. Each of these categories of products does so using different mechanisms.

MDM

Mobile Device Management (MDM) products aim to provide information security for devices that live outside of the organization's perimeter. Cell phones, tablets, laptops, are all targets for MDM. MDM products apply security templates, profiles and policies to remote devices, while ensuring that mobile devices meet organizational information security requirements before they are allowed to connect to resources behind the corporate perimeter with a VPN or Remote Access technology.

The secure delivery of applications is another consideration for MDM, with solutions to this problem ranging from app stores to Virtual Desktop Infrastructure (VDI). MDM products also focus on access control issues: ensuring that only the authorized user can use the device, and that the device can be tracked or remotely wiped if stolen or lost.

AUTHENTICATION

Centralized authentication and Unified Authentication (UA) technologies are what they say on the tin. Modern centralized authentication relies on directory services such as LDAP, SAML, or Microsoft's Active Directory. Single Sign On (SSO) is the most recognizable UA technology, and aims to allow users to use a single username and password to access workloads and services from multiple providers, located on multiple different infrastructures.

UA is more than SSO, however. UA technologies also need to consider how access control will translate between the different authentication structures, and how they differ between technologies. The major directory services are similar enough, but SSO integrations with popular Software as a Service (SaaS) applications can introduce interesting wrinkles.

Multi-Factor Authentication (MFA) systems can also cause complications, especially if they need to differ based on country. Consider SMS verification, a common MFA approach is requiring users to enter a code from a text message to log in. Many countries implement restrictions on SMS messages that can make this unreliable, or at the very least require additional integrations to work properly.

AUTOMATION

The most important information security defense of them all is automation. Service provider networks have long been impossible to manage without automation. Enterprise networks are at this point today. Midsize and small business networks are beginning to see this problem affect them, and it's only going to get worse.

The number of workloads under management is not only growing, that growth is accelerating. The diversity of workloads is constantly growing, and the rate of change in IT hasn't stopped accelerating for some time. Humans just can't keep up without help, and this is what automation is for.

Due to the highly interconnected nature of today's IT, the rapidly vanishing network perimeter, and the fact that infosec threats can and do happen from behind the perimeter, infosec must be a part of all aspects of IT. Every piece of IT infrastructure, every workload, and every network-connected device must be part of an organization's infosec design.

Because of this, IT automation is infosec automation. The same forces that create a requirement to automate IT also drive the need to automate security. The bad guys are going to use as much automation as possible in attacking, and their greatest weapon is also the defender's most potent shield.

04

ANATOMY OF AN ATTACK

Having just read about a big list of information security technologies, it's probably easy to assume that any hacker (or defender) has to be a genius nerd wizard who knows every single possible exploit and how it is usually countered. That's only marginally true.

Defenders don't have to write their own security applications; they can buy products and services that do most of the heavy lifting. Similarly, cybercriminals don't have to write their own malware. There are thriving dark web commerce sites where they can buy someone else's exploitation tools and instructions for their use.

Both attackers and defenders need to have enough technical skill to configure the products, test them, and make them work together. But as discussed in chapter one, information security is not about specific technologies so much as it is about a way of thinking.

Hackers often learn to think like their targets so that they can pull off a piece of social engineering. Infosec practitioners learn to think like hackers so that they can figure out what weak points to defend. Remember, the weak parts are not usually the computers. The weak parts are the people. Information security is mostly about mindset.

To give some insight into the mind of a hacker, let's look at a hypothetical example of how a successful hack might occur. Readers should note that, while this may seem detailed, it is in fact a very high-level overview of the anatomy of a hack.

In the real world, there are even more steps required to successfully complete a hack, which are not going to be explained here for obvious reasons. If you attempt to use this as a guide to hack, you will fail. If you attempt to use this as an information shopping list to go careening around the dark web, looking for answers to fill in the gaps, you will get vanned.

For those unfamiliar with the term, to get vanned is to be arrested by law enforcement agencies. This generally considered an unpleasant experience. Do not to attempt this at home, or anywhere else.

More succinctly: don't do crimes. It doesn't end well.

DEFINING THE OBJECTIVE

Imagine a hacker; we'll call him Bob.

A side note on gender: it's not that there aren't female hackers – there most definitely are – but for whatever reason, there are substantially more male than female hackers. Also, while "hacker" is a broad term that can include security researchers, programming enthusiasts, and other non-criminal nerds, the subset of hackers who commit crimes ("black hat" hackers) is even more likely to be male than their law-abiding counterparts. Bob the black-hat hacker is male for statistical reasons.

Say that Bob's target is PotatoCom, a local internet service provider. Bob wants to compromise the database underpinning PotatoCom's financial system. Since mindset is important, we need to know why. Let's say that instead of profit, our hacker wants revenge. PotatoCom has refused to install fiber-optic internet for the umpteenth year in a row. Bob is sick and tired of terrible ADSL.

If this seems like a trivial reason, bear in mind that hackers don't always wreak havoc for profit or as part of an ominous state-sponsored espionage plot. They might just as easily hack for fun, to show off, to correct what they see as an injustice, or for revenge. The results are still unpleasant.


.COM

Because revenge is Bob's goal, he doesn't want to actually use the contents of PotatoCom's database, but he wants to deny PotatoCom access to it. This affects how he designs and executes his attack.

First, he has to ask himself, "How do I deny PotatoCom access to their database?" Databases tend to be both clustered and backed up, so simply disabling an individual host won't do it. Bob needs to do damage in a way that will replicate to the other database servers in the cluster. He also needs to ensure that the backups are damaged, otherwise PotatoCom can just restore from a backup and keep using their financial database.

To really make PotatoCom hurt, our hacker would therefore have to scramble the data in the database. This means corrupting records in the database rather than simply encrypting the database files, or the host operating system.

Now the hacker needs a way to gain control of the relevant system(s) so that the attacks can be accomplished. To do so, Bob doesn't need to be some super-technologist who knows every possible exploit of every product. He only needs to learn enough to achieve his goals. (Remember, he can purchase malware on the dark web.) To find out the rest of what he needs to know, Bob is going to do some reconnaissance.

He has three reconnaissance targets: identify the database that PotatoCom uses, identify their backup mechanism(s), and identify a way in.


.COM

CONNECTIVITY

Before proceeding any further, we need to take a detour to discuss connectivity. A lot of the parts to Bob's process will involve going online. In order to not get caught, he needs to find a way to perform his illicit online activities in a way that cannot be traced to him. There are a few different approaches available.

One such approach is using unsecured public Wi-Fi, which might be available from a fast food outlet or coffee shop. This has two hurdles to overcome. First, Bob needs to be able to access the Wi-Fi without getting caught. So he needs to make certain not to show up on surveillance cameras.

Finding a place without any cameras would be ideal, but this is becoming increasingly difficult. He might find an unsurveilled place to connect to the Wi-Fi, but he also needs to think about cameras he might pass while getting there. To protect himself in case he passes a camera he cannot avoid (or failed to detect), he might use physical disguises. He will never use the same Wi-Fi hotspot twice.

Ideally, Bob will vary the type of transportation he uses get to his various free Wi-Fi locations. He's trying to avoid creating a noticeable pattern of activity. He might use a bus, or a taxi service that will accept payment in cash (so long as he has never handled the cash with his bare hands, so it is clean of his fingerprints.) He could park someplace nearby and walk the last few blocks.

He might not have to enter the building of the business providing the free Wi-Fi, if he can get close enough to connect. He could work from an adjacent building or a public park, if he can find a spot without surveillance.

This solves the problem of getting online in the first place. The other hurdle Bob faces is the need to hide his activities from any surveillance technologies that exist on the networks that he is traversing. He will probably need to use a VPN, or more than one. He might also rent a VM or Virtual Private Server as a place to stage his attacks. To further hide his internet traffic, he will probably use TOR, I2P, or another method of anonymization and encryption.

HERE'S ONE POSSIBLE PATH BOB MIGHT TAKE TO STAY AS ANONYMOUS AS POSSIBLE:

1. Bob creates a Linux Live CD that is set up to automatically connect to a paid-for VPN service. This image would also be configured to randomize the MAC address of all network devices. (While devices ship with a hardware-encoded MAC address, computers can be told to lie about their MAC addresses with ease.)

2. He saves this image somewhere on an anonymous cloud service to be accessed on as-needed basis.

3. He burns the Live CD to a single-use USB stick for each "hacking" operation. When he plugs the USB stick into any computer, he can boot from it instead of using the computer's pre-installed operating system. The USB stick will automatically connect him to a VPN service.

4. Bob uses "disposable" notebook computers that have nothing on them, ones that he would be perfectly comfortable ditching if he needed to do so. He can leave them with their default install, since he will be booting off the USB stick anyway.

5. He connects to the VPN service immediately after connecting to the free public Wi-Fi.

6 Then Bob connects to TOR over the VPN service.

7 He passes an SSH tunnel through TOR to the first of the private VMs he has rented. This lets him use that VM in all the same ways as if he were sitting in front of a real, physical computer.

8 He gets the private VM to connect to I2P. Most defenders block TOR/I2P exit nodes, so he can't go straight from I2P to connect to the things he wants to hack.

9 Instead, he goes through I2P to connect to a second private VM.

10 From there, Bob connects to a second (different, classier) paid-for VPN service.

By now, Bob hopes he has thoroughly muddled his tracks. He comes out the other end of this process looking like some relatively harmless internet user, with no obvious connection back to the black-hat hacker using the public Wi-Fi at a fast food restaurant.

This is a lot of steps, and this isn't everything Bob needs to do to pull off the hack – not by a long shot. But he needs to do these steps every single time he connects to the internet during the hack against PotatoCom, or he risks getting caught.

PAYING

You'll note that Bob has a few problems in the scenario above, and they're all money-related. To do this set of steps every time he wants to go online, Bob needs to buy:

1. USB Flash drives
2. Notebook computers
3. Two VPN service accounts
4. Two private VM rentals

At least one of these requires him to pay in currency that can be tracked by law enforcement agencies: the classy VPN, and probably the private VMs too. So he has to think through how to do that without being identified.

Today, it is possible to get pre-paid credit cards, though these increasingly require people to give up personal information to get them, and it's pretty much guaranteed that anyone who does get one will be on camera doing so. Bob needs a way to do this without being traced.

The most popular ways to buy untraceable pre-paid cards are to buy them in bitcoin, or to trade gift cards from major retailers for them. As with the ways Bob gets online without being detected, the key here is being indirect. He needs not to leave a trail, or to only leave one that is tangled and difficult to trace. Basically, he needs to reverse-launder money.

If Bob can find a bitcoin ATM that provides bitcoin without taking his picture, and there's no surveillance on the entrance and exit routes, he's in luck. Assuming the bitcoin ATM takes cash. (Most don't.) Let's assume he doesn't get lucky, so he has to to give up on bitcoin (which is traceable anyways), and focus on gift cards. For hackers, gift cards are fantastic.

Every year, just after Christmas, the classified sections all over the world fill up with people trying to get rid of gift cards. A small amount of research will allow Bob to find gift cards that have minimal tracking, and there are entire forums dedicated to telling people which are the best gift cards for exchange purposes. Minimal tracking, for Bob's purposes, means the gift card vendor doesn't take any personal info or track usage.

Bob isn't going to buy those gift cards straight away. He's still adding layers of misdirection. Bob will go buy gift cards from stores people actually want to buy things from (Amazon, Apple, etc.). He will buy a few, never buying more than one from the same place.

These can almost always be bought in retail and grocery stores. Again, Bob will probably use simple disguises to conceal his identity. Just looking tired will prevent cashiers from asking him too many questions. If he is extra paranoid, he will buy from a grocery store he never frequents when he is actually buying groceries.

Maybe he will do this while he is travelling somewhere. Bob can mail the gift cards back to himself without them getting inspected if he doesn't mail them through customs. Or better yet, he can mail them to a neighbor nearby that he knows is at work during the day. Bob can put a small Wi-Fi camera on a nearby tree, and simply watch and wait until the delivery comes. Then he can pilfer the delivery. (In disguise, of course; his neighbor might have a surveillance camera.)

Now Bob goes back online to exchange these gift cards with people who have gift cards that have low tracking. He should be able to swap 1:1 here, so he doesn't lose money (yet). Then (still online) he swaps his new, low-tracking gift cards for pre-paid credit cards. (He probably only gets about 70% return on these swaps.)

To recap: what Bob has just done is taken cash which is traceable to him, converted it into a gift card that most people actually want to spend, and traded it for gift cards that are less in demand, but where the gift card vendor doesn't take any personal info or track usage. Those cards are then traded for pre-paid credit cards, giving Bob money that he can use anywhere, and which cannot be traced back to him.

These pre-paid credit cards can be used to buy things like VPN service access, private VMs, USB flash drives, and notebooks. If Bob is being very careful, he will never use the same card to pay for multiple services or devices. He will also try very hard not to get multiple cards from the same person, or even from the same forum.

RECONNAISSANCE

Now that Bob has untraceable money, and has a way to connect to the internet without getting caught, let's turn our attention to the hack itself. Remember, Bob needs to find out three things: which database PotatoCom uses, what their backup mechanism(s) are, and how to get into their systems.

To make this explanation simple, we're going to assume that PotatoCom uses Amazon's AWS for everything. Bob can easily figure this out through a number of different means. He could use technical means to probe PotatoCom's IT infrastructure. Alternately, he could use social engineering to chat up various PotatoCom employees, presenting himself as a knowledgeable individual (or even an insider) and giving people scope to vent. This usually works.

Bob could also find what he needs on social media. All he needs to do is look for PotatoCom employees griping about the tools they use internally to their organization.

He could even try calling up some of the vendor support lines and pretending to be from PotatoCom. He doesn't need to make changes to anything, just confirm that PotatoCom is a customer.

Once Bob has identified that Amazon is the host of PotatoCom's infrastructure, there are several options to complete his reconnaissance. The simplest method is to socially engineer someone into giving him valid credentials for login. Administrative credentials are emphatically not required for this step; anyone who can "see" the infrastructure is just fine, even if they can't modify it.

Let's assume that, through whichever method, Bob figures out that a) PotatoCom's infrastructure is located on AWS, b) that they're using an Amazon-provided MSSQL cluster, and c) that they're using Amazon's native backup capability.

There are a few ways in. One would be to try to crack one of the workloads running on AWS. Bob would then use this toehold into PotatoCom's infrastructure to go after something bigger. He might start with some network scans to probe the infrastructure, identify what's there, then attack soft targets. (This is referred to as "lateral movement".)

Once there are multiple backdoors into the system, Bob could try super tricky technical stuff to get into the database and corrupt entries. Maybe some voodoo could even be performed to convince Amazon not to back up the databases, or to jettison existing backup copies.

This is the hard way. Bob would probably have to spend months researching how to do it, because compromising each application and operation system is its own challenge.

He would probably have to build a replica of PotatoCom's environment, use default settings, then see if he could break in. He would then build another replica environment, use the most loudly published "industry best practices secured configuration" version of that same deployment, and attack it until he could break that. At this point, Bob would finally have an exploit to try on PotatoCom's infrastructure.

But that's an awful lot of unnecessary work. All Bob really needs to defeat PotatoCom's security is their administrative password to Amazon. They've got their entire infrastructure on a single cloud provider, so it's a good bet that they use the same administrative credentials for their production environment as they do for their backups. All Bob has to do is get those. That's way easier.

Of course, if PotatoCom were more security-savvy, they wouldn't be using a single account for production and backups, or even a single cloud provider. But, hey, it's PotatoCom. Their apathy, poor quality of service, and sloppy business practices are what made Bob angry at them in the first place. (Also: most organizations are this dumb. Honestly.)

GETTING TO KNOW THE TARGET

Bob knows what he wants to accomplish, and has a way to pull it off without getting caught. He just needs administrative credentials.

The easiest way to get credentials is to socially engineer someone to give them to you. That might work for someone with low-level access, but Bob is unlikely to be able to talk someone with administrative credentials out of theirs. (They will be a lot more aware of the possibility of credentials being misused, and will be familiar with common phishing techniques.) So he turns to physical surveillance.

Again, this is a multi-step process. You may have noticed by now that hackers need to be patient, methodical, and creative.

1. Bob needs to identify PotatoCom's cloud administrator(s). Who are they?

2. He has to learn the routine of PotatoCom's administrator(s). He uses this to figure out his next steps.

3. Does the target admin work from home some or all of the time? If so, Bob is in luck. Homes are never secured as well as offices, and there are LOTS of opportunities to spy on people in their own home.

4. If the target admin doesn't work from home, where is their office? If it's by a window, and the desk is laid out right, Bob might be able to get their login credentials with a simple telescope. If the layout is not convenient, he might have to smuggle in a camera somehow so that he can get a look at their keyboard.

5

Bob could learn the target admin's hobbies. Social media is useful here. If they like a particular science fiction show, humorous paraphernalia from that show could entice them. The camera could be hidden on or in the kitsch.

6

Bob might also consider taking a temp job with the cleaning service that cleans PotatoCom's offices. This would give him lots of opportunity to plant a camera. A cleaner has a plausible reason to be wearing rubber gloves, so he can avoid leaving fingerprints. If he takes this route, he won't go forward with the rest of the hack until well after he has quit the cleaning job.

7

Alternately, the target could be using a wireless keyboard. Some of these can be snooped on with simple passive electronic sensors.

8

A keylogger would also work, if Bob cannot place a camera. Bob's best bet would be to try to get one that transmits its data wirelessly. He can stick an inconspicuous booster on the side of PotatoCom's building to pick up the signal. The booster will send the signal to hardware that Bob can stash somewhere better concealed. This hardware will talk to the internet via a burner phone operating in hotspot mode, or it might even use PotatoCom's own free guest Wi-Fi.

9

Spear Phishing is another possible tactic. The infamous lost USB drive trick is an example of an attack that shouldn't work against any competent sysadmin, but sadly works distressingly frequently.

There are more possibilities, but you probably get the idea. If Bob learns about his target, he'll find the vulnerability that lets him get a camera where it needs to be. Thus he gets the administrative credentials. Finally, he's ready to take his revenge.

He connects to PotatoCom's AWS infrastructure via his untraceable tunnel bought with untraceable money. He disables any change reporting features that exist so that the changes he is about to make aren't reported to PotatoCom's IT team. (His camera surveillance will have alerted him to the existence of change reporting, if PotatoCom uses any, and probably gained him any extra credentials he needs in order to defeat it.)

He corrupts the database. He then performs a backup or two so that the latest backups have the corrupted data. He deletes the older, uncorrupted backups, changes PotatoCom's AWS password, and burns any evidence.

PotatoCom's financial database is now basically unusable, and they can't get access to their AWS account to figure out what went wrong. They cannot bill their customers, pay suppliers, or file their taxes. They might trigger an audit and then get in trouble for not having access to the required data. They might even go out of business.

THINKING LIKE A HACKER

What you should take away from that mental walkthrough is that nothing the hacker did here was technically difficult. Even the parts that might have seemed like they required advanced technical skill would have been easy to learn. Many of the individual pieces of Bob's process are totally legal, and the internet provides free, step-by-step tutorials on how to accomplish them.

However, it was the non-technical approaches that were the path of least resistance in almost every case. What mattered to Bob wasn't a fist full of zero-days, arcane mastery of the Linux command line, or in-depth knowledge of individual security products.

What mattered was mindset.

First: never get caught. Bob went to great lengths to make it hard to identify him. You'll notice that the actual hack took only a few sentences to describe, while the procedures Bob followed to keep himself anonymous took up most of his time (and yours). Hackers are patient, creative, and usually paranoid.

An addendum to this first rule of hacking is never, ever reuse an identity. Reusing identities gets you caught. The expansion of this point is that hacking is expensive, because creating new, untraceable identities is expensive.

Second: always look for the exploit. There is no point in taking the hard way if an alternative is available. The most exploitable part of any system is usually the human being responsible for it.

That's how a hacker thinks. And to defeat a hacker, a defender needs only the right mindset: safety first.

Which security tools a defender uses will vary greatly from area of responsibility to area of responsibility, even within the same organization. The defender will have to learn how each of the applications, operating systems, network devices, and cloud providers that they are responsible for work. Then they can start to think like a hacker - looking for the exploits - and plug those gaps with technology, business processes, or even just simple curtains.

Defenders discover the tools they need to use once they know what they are defending. Just as with hackers, the mindset is what matters.

05

ADVANCED CONCEPTS

When talking about information security technologies, much emphasis is put on prevention. As the old saying goes, “an ounce of prevention is worth a pound of cure”. But prevention alone isn’t enough.

Detection is an even more important concept than prevention, and worth a significant investment. It doesn’t matter who you are, or whom your organization is, you will be compromised. Detection technologies – such as monitoring, SIEM and ATP products – let organizations know when they’ve been compromised.

Information security compromise events are simply a part of life. This isn’t an attempt to be flippant. Every year statistics are released that show a significant number of workloads, systems, and devices are compromised without the knowledge of their host organizations.

The inevitability of compromise doesn’t invalidate prevention technologies, but it does mean that organizations must plan for what happens when a compromise occurs. In addition to investing in detection technologies, organizations need to invest in mitigation technologies – NGAV, UA, browser defenses and so forth – as well as implement IT policies that aim to stop the lateral spread of threats.

Lateral compromise occurs when a workload or device located behind the organization’s primary network defense perimeter is compromised, and then the attacker takes advantage of the more lax security behind the perimeter to compromise other workloads or devices. Lateral compromise can allow a single compromise event to cascade, making incident response planning a modern infosec necessity.

Incident response planning should lean heavily on automation. By tying automation into detection and mitigation technologies, automated incident response can be achieved. When a compromise event is detected, automation can be used to trigger mitigation steps, such as ordering a network switch to disconnect or quarantine a compromised device. This prevents lateral spread, and can save enterprises tens or even hundreds of millions of dollars per event.

REGULATORY COMPLIANCE

A basic understanding of infosec concepts, technologies, and techniques is a requirement of a number of infosec-related regulatory compliance standards. Increasingly, it is not enough to simply implement information security technologies, or pass an audit by applying IT policies in a tick box fashion. Organizations are being asked to demonstrate that they actually understand the basics, and why information security is necessary.

The European Union's (EU's) General Data Protection Regulation (GDPR), for example, requires any organization involved in large-scale data handling to have a Data Protection Officer (DPO). DPOs are responsible for guaranteeing organizational compliance with the GDPR, and there are personal ramifications if the DPO or organization's they work for are negligent in securing the data of EU citizens.

The GDPR seeking to ensure that there are personal consequences for an organization actions is currently the most notable legal effort to pierce the corporate veil, but it is not unique. Prominent politicians in all major Western nations have been openly discussing exactly this for years as tech giants like Google and Facebook fall out of public and political favor.

The never-ending drumbeat of high profile data breaches by organizations and governments has only added fuel to the fire. Infosec, privacy, and regulatory compliance are inextricably intertwined.

NETWORK SEGMENTATION AND MICROSEGMENTATION

While the traditional organizational network perimeter may be disappearing, the concept of a network edge is not a bad thing. Networks are broken up in several ways. The traditional network defense perimeter was built around the router that provided an organization access to the internet, and is a legacy from an era where most organizations only had one internet connection.

Even today, many large, distributed organizations use services such as Multiprotocol Label Switching (MPLS) and Virtual Private LAN Services (VPLS) to interconnect locations, funneling all internet-bound traffic from all locations through a single network edge. Traditionally this was done because the cost of both internet connectivity and infosec products were prohibitive. This is no longer the case.

Today it is entirely common for multiple locations to have internet connectivity, and for each to be defended by its own perimeter defenses. The cost of both internet connectivity and infosec technologies had to drop to get there, and today's price points are enabling even more innovation.

Network segmentation is the practical form of an acceptance that threats can come from anywhere. Implementing network segmentation means that in addition to traditional "north-south" defenses, network administrators are preparing to guard against compromise from deep within the network. These "east-west" defenses are geared towards preventing lateral movement of attackers who have succeeded in compromising a system on the network.

Subnets, VLANs, and VXLANs, and L3VPNs can all be used to isolate one group of devices from another. Allowing interconnectivity between these network segments requires a router, making routers a natural chokepoint for traffic inspection. Additionally, ACLs can be applied to each point of connection along a network.

The term microsegmentation hypothetically means the practice of deliberately breaking up networks into small segments. The precise definition of microsegmentation depends on whom you ask. Broadly, microsegmentation means breaking the network into a bunch of very small segments, typically smaller than would ordinarily be achieved using VLANs.

The nomenclature isn't settled, however, as some vendors and communities believe that unless a specific technology or approach is used, it doesn't "count" as microsegmentation, and is merely traditional network segmentation. Regardless of how you define it, the basic principle is the same in each case: enforce network policy as close to the system to be defended as possible.

In most microsegmentation designs, each network segment contains a single application, with each segment being defended by a router, firewall, and other security tools. Ideally, however, every point of connection on a network is defended, an approach that is increasingly important given the sprawling size of modern applications. Today's applications can consist of dozens, hundreds, or even thousands of individual workloads and devices, and can be spread across multiple public clouds, edge computing data centers, as well as on premises.

Software defined networking and automation are essential to enable microsegmentation, especially when distributed, cloud-native, and/or microservices-based applications are in use. Each microsegment has its own network edge, and just like the traditional internet-facing network perimeter, microsegment network edges can be defended.

Ring fencing individual applications with their own layer of network-based defenses allows those applications to be protected against lateral attack. It can also mean that those applications can be more readily made available to employees or customers who wish to access them over the internet without the use of VPNs or Remote Access technologies.

Applying a network wrapper for individual applications is ambitious. Applying similar levels of network defenses – or even simply constraints in the form of ACLs – to every single connection on a network is even more so, but these are not the only use cases for breaking up networks. Network segmentation (including microsegmentation) is also useful in shared environments, where organizations wish to build (or participate in) IT infrastructures that have multiple tenants. Multitenancy is a critical concept for cloud service providers, and edge computing deployments.

A real world example of microsegmentation in use is a research hospital. Medical records have strict regulatory compliance concerns, and patient data must be fiercely guarded. Research data has special restrictions placed upon it, and doubly so if government money is involved.

The traditional approach to stringent regulatory compliance requirements would be to have each research project purchase its own IT infrastructure. This is economically inefficient, not to mention time consuming.

Private clouds, combined with microsegmentation and appropriate infosec measures can completely isolate a research project as tenant on the shared hospital IT infrastructure. In this sort of multitenant environment, each tenant would be contained in their own microsegment, with the virtual network edge defended by a full suite of network services, just as if they had dedicated physical IT equipment.

THE INTERNET OF 5G THINGS

As we head into the 2020s, we are seeing two major trends in IT collide. Internet of Things (IoT) devices – autonomous devices designed to connect to a network and which typically operate without human interaction – have reached critical mass. Their adoption rate has been accelerating for years, but we are closing out the 2010s with 2019 being projected to see the number of IoT devices surpass 25 billion, with a total global market estimated at $1.7T.

The other major trend leading into the 2020s is 5G connectivity. 5G is the next generation mobile/cellular network technology, designed to allow lots and lots of IoT "things" to connect to it. This is being enabled in large by certification of Mobile Network Operator (MNO) operated picocells. These picocells operate at 60 GHz, meaning that the radio signals won't leave an individual room, and so reside inside an organization's premises.

Unlike Wi-Fi, which is typically owned, serviced by, and connected to the organization's own network, picocells are part of the MNO network. 5G picocells are not merely a way to get better reception on one's cell phone indoors; they are designed to have large numbers of IoT devices connected to them.

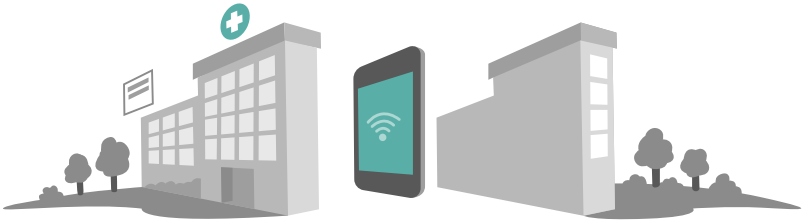

A hospital, to reuse an example, can have thousands – tens or even hundreds of thousands – of IoT devices in a single building, even today. The number of IoT devices in the most connected hospitals is astonishing, and their use is exploding.

Hospitals are a great example of how the entire suite of infosec technologies comes into play. None of us can afford to have an emergency room's IoT starved for bandwidth because the printers in accounting were compromised after someone watched a booby-trapped cat video. Service providers need to be able to microsegment their 5G networks, and defend each segment with a wide array of information security technologies - and they have to do this at incomprehensible scale.

The hospital's IoT devices will have to communicate with one another. They will have to communicate with the hospital's own network which, while located in the same building, is a completely separate network that is "across the internet" from the 5G network provided by the MNO. The IoT devices will likely also have to communicate with applications and services hosted in public clouds and edge computing datacenters, and all of this has to happen in a regulatory compliant fashion.

PUTTING IT ALL TOGETHER

Firewalls, including WAFs and NGFWs, as well as anti-malware and NGAV products are only the beginning of information security. Encryption, monitoring, DLP, CASBs, SIEMs, ATPs, VPNs, Remote Access, Browser Defenses, MDM, UA, Automation and microsegmentation are all important tools in the information security professional toolbox.

Prevention, detection, mitigation and incident response are the four areas of infosec responsibility to which organizations must attend. But as you explore and consider all of these various technologies, remember that information security fundamentally isn't about IT, or any technology at all. Information security is a mindset, one that even our friend, the amiable but simple nuthatch can master.

Humans are an organization's greatest information security vulnerability. But with the right mindset – one focused on safety – they can be its greatest asset as well.

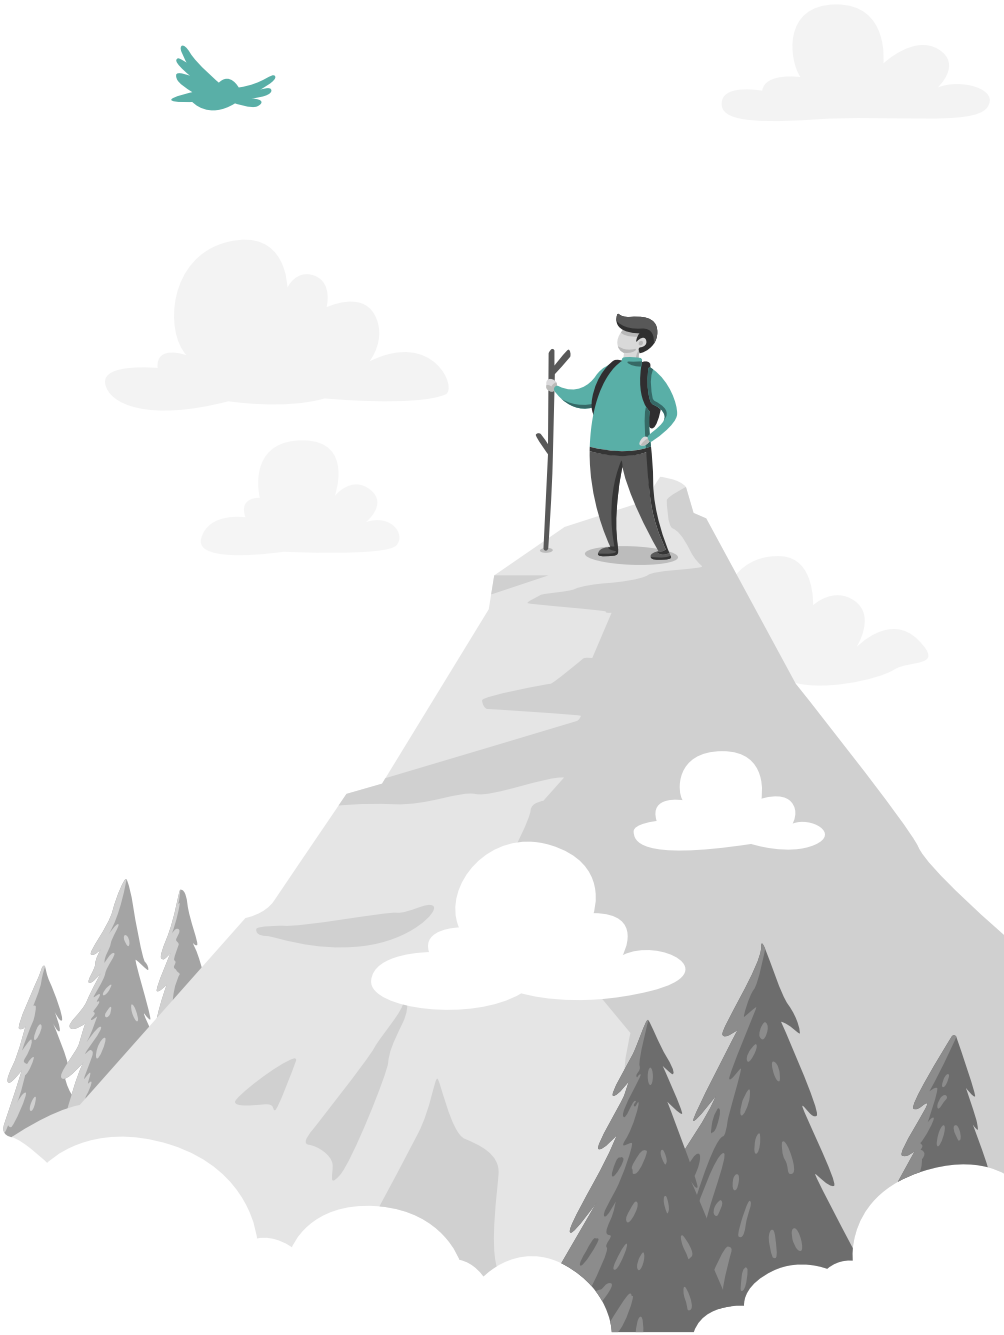

06

JUNIPER CONNECTED SECURITY

An attack can occur against any part of an organization's network. The point of initial compromise could be an IoT lightbulb, someone's Windows computer, or even an improperly configured network switch. Anything on a network – including the infrastructure that makes up that network – can be targeted. Defending today's networks means defending everything on those networks, and that requires a different approach to information security than is practiced by most organizations.

Two things are required to defend a network: visibility into threats, and a means to counteract those threats. Network visibility comes from telemetry. Log files, live streaming data feeds, and other forms of monitoring are all telemetry that can be fed into centralized systems like SIEMs or ATPs. These tools perform analysis on the telemetry they receive, and present that data in a form that can be acted upon, either by a human or an orchestration platform.

Where vendors attempt to differentiate themselves is generally in the means they offer to counteract threats. The introduction to information security concepts above briefly – and rather broadly – introduces the main product categories used by information security vendors today, however, even the umpteen pages that you have had to read to get this far only scratch the surface of what is possible, to say nothing of the myriad ways in which individual vendors' offerings differ.

SECURITY NEEDS LAYERS

Juniper Networks believe that networking and security are inextricably intertwined, and that it is when vendors (and customers) attempt to architect their networks using point solutions that things tend go horribly wrong. Juniper believes that effective network security comes from interconnecting multiple layers of security, and that layered security is achieved by deploying multiple technologies, from multiple vendors, to achieve a whole that is more capable than the sum of its parts.

Switches, routers, and Wi-Fi access points participate in Juniper Connected Security as both providers of deep network visibility and points of network policy enforcement, all centrally automated and orchestrated. This provides organizations the east-west protection that is often so difficult to achieve in networks composed entirely of point solutions.

The “best” solution (according to analyst rankings) for a given segment is likely to be one provided by a specialist point player. A point solution provider’s entire company is focused on the creation of that one product, and they will inevitably expand the features and functionality of that product beyond what is traditional for the category, because they have exactly one card to play, and achieving usable network defenses requires more than any one product can possibly provide.

In order to maintain their leadership position, the point solution vendor needs to keep adding features, however; this approach has limits, and in order to secure their networks, organizations will have to use multiple products. If the organization is one which only buys products with the highest analyst ranking, then the hard work of securing that network will be in getting all the various bits to play nice with one another.

This approach can also make automation and orchestration of network defenses difficult: rankings change all the time, and automation implementations frequently outlast the lifecycles of the individual products they are automating. Automation is what makes modern IT possible, and nearly every organization today is utterly dependent upon IT. The harder maintaining IT automation becomes, the harder it is for organizations to operate at all.

There are multiple vendors which have a diverse portfolio of networking and information security products, including partner ecosystems. There are also multiple vendors which sell to customers operating at service provider scale, and which as a result we can handle the craziest, most outsized networks that any enterprise customer can dream up.

What sets Juniper apart is the “Connected” in “Juniper Connected Security”, specifically, our commitment to interconnectivity. Juniper champions integration and orchestration. Juniper encourages the use of open standards, open protocols, open APIs, and even builds in support for competing products: Juniper’s goal is to help customers make the best use of what they already have, instead of demanding that they rip and replace.

Juniper's commitment to using JunOS everywhere is a part of this strategy, and a further point of differentiation. Juniper believes that reliability is the most important metric for any network, and that automation is the enabler of network reliability.

Network reliability engineers need a multicloud-aware, feature-rich, and common management stack upon which they can build the automation their organizations will rely. This management stack needs to be extensible so that organizations can take advantage of emerging technologies in their quest to realize the benefits of becoming AI-driven enterprises.

The use of a single operating system, JunOS, throughout Juniper's portfolio makes centralized management of Juniper products much simpler, and enables Juniper to offer feature rich management platforms with management plans that are on-premises, cloud based, or a combination of both. It is this common operating system that makes adding functionality throughout the portfolio economical. Instead of being an agglomeration of point solutions, each effectively operating as their own entity, Juniper Networks regularly extends security functionality throughout the network.

Juniper Connected Security gives organizations the ability to safeguard users, applications, and infrastructure by extending security to all points of connection across the network. By enforcing policy as close to a threat as possible, network reliability engineers reduce the risk of that threat spreading. Through the use of machine learning, advanced analytics, and automation, rapid incident response becomes a reality.

PN 7400117