TECH REVIEW 2023

NETWORK BOX

www.network-box.com

TECH REVIEW 2023

Table of contents

As a special end-of-year review, Network Box has complied the key *In the Boxing Ring* technology news, features, and articles from 2023.

To read all available editions of *In the Boxing Ring,* please use the link provided below:

https://www.network-box.com/itbr

www.network-box.com

Network Box Services in 2023

2022 was a tough year for many. Putting COVID-19 in our rearview mirror, we are trying to deal with managing the hybrid networks that have resulted from three years of Pandemic work-from-home and the proliferation of mobile devices and Software-As-A-Service deployments. Combine that with budget restrictions in the economic downturn, and things have been challenging.

But despite so much changing, the fundamentals remain the same. Despite dropping from its peak in 2021, ransomware continues to be the #1 malware problem, with data breaches topping the list of threats. Still, 80% of security incidents are caused by a lack of protection, and the other 20% by that protection not being configured/maintained correctly. Network Box's approach of combining a full suite of protection technologies with PUSH protection updates and configurations managed by our engineers in Security Operation Centres around the world has continued to prove highly effective in keeping the networks we manage secure, addressing both the 80% and 20%.

2023 Roadmap

The year saw Network Box continue to expand on our core technologies, broadening our service offerings and evolving to improve our Managed Detect and Response service. This means:

- NBSIEM+ and our mobile apps were expanded to form the admin and user portals to all our services unified under one umbrella.
- A dashboard facility was released for NBSIEM+ to allow users to design their own dashboards (in a similar way to the admin web portal). More reporting and KPI widgets will be introduced to support this.
- Improved reporting, as well as providing clear advice to our customers regarding their information security, were key points of focus this year. We introduced two new major types of reports:
    1. Conducting regular configuration reviews of each Network Box under management against the 16 categories and 81 items of our Best Practices; to periodically report on these and review areas of concern with our customers to address identified shortfalls. These Best Practices represent the most common forms of network infiltration and data breaches affecting networks worldwide. Network Box Security Engineers refer to these when designing defense systems for networks under management, when processing policy change requests, and during such periodic configuration reviews.
    2. A new Executive Summary Report produced by NBSIEM+ will periodically summarize the KPIs and services provided at a high level, which is able to report across all devices and cloud services provided in one single report, organized by configurable asset hierarchies. This report leverages Network Box's unique capability of storing data locally (on-end protection devices) but being able to report centrally in a consolidated manner while adhering to regulations such as GDPR.
- Expansion of our services to the desktop. This year, we started offering an end-point protection service, as well as furthering our existing vendor partnerships with improved NBSIEM+ support. While the gateway protection provided by Network Box appliances (physical, virtual, and multi-tenanted cloud) is comprehensive, end-point protection provides for in-depth defense improvements and support for devices not within or traveling outside the protected networks.
- We continued to enhance and expand our cloud service offerings - heading towards NBSIEM+ as the single unified platform.
- Finally, we continued our work on our new NBRS-8 platform and commit to offering upgrades to this platform on all current Network Box hardware when it is released. More information on this will be released soon.

2023 saw new challenges, and just as we have over the past 20+ years, Network Box rose to meet these. As a managed service, we are continually updating and adjusting our offerings to meet the daily threats facing our customers. We would like to thank all of you for your continued trust in Network Box, our platform, and our security services.

Understanding Company's SECURITY POSTURE

As businesses increasingly rely on digital infrastructure and data, ensuring a robust security posture is paramount. In this article, we will delve into the concept of a company's security posture and explore its significance in today's threat landscape.

Organizations face various security risks in the digital age, including cyberattacks, data breaches, and insider threats. A company's security posture refers to its overall approach and readiness to protect its assets, systems, and sensitive information from these risks. It encompasses various elements, such as policies, procedures, technologies, and employee awareness.

Critical Components of a Strong Security Posture

Risk Assessment: Conducting a thorough risk assessment helps identify potential vulnerabilities and assess the likelihood and impact of various threats. This forms the foundation for developing effective security measures.

Security Policies and Procedures: Establishing comprehensive security policies and procedures ensures consistent adherence to best practices and regulatory requirements. This includes defining access controls, incident response protocols, and data protection guidelines.

Technology Infrastructure: Implementing robust security technologies, such as firewalls, intrusion detection systems, and encryption tools, fortifies an organization's defense against external and internal threats. Regular updates and patch management are essential to address emerging vulnerabilities.

Employee Awareness and Training: Human factors are critical in maintaining a solid security posture. Educating employees about cybersecurity best practices, phishing awareness, and social engineering tactics helps foster a security-conscious culture.

Incident Response and Business Continuity: Preparing for security incidents and having a well-defined incident response plan is crucial. This includes establishing roles and responsibilities, conducting regular drills, and having backup and recovery strategies to minimize disruptions.

Continuous Improvement and Adaptability

Maintaining a strong security posture is an ongoing process. Organizations must stay vigilant, adapt to evolving threats, and continuously improve security measures. Regular audits, vulnerability assessments, and penetration testing help identify weaknesses and enhance defenses.

A company's security posture is critical to its overall risk management strategy. Organizations can protect their assets, maintain customer trust, and ensure business continuity by prioritizing security. A collective shared responsibility requires collaboration between IT teams, employees, and stakeholders to create a secure environment.

Shared Responsibility in Cybersecurity

Cybersecurity is a shared responsibility in today's interconnected world. It is no longer solely the responsibility of organizations or individuals to protect themselves from cyber threats. Instead, it requires a collaborative effort between all stakeholders to ensure a secure digital environment.

Organizations must play a vital role in cybersecurity by implementing robust security measures and protocols. They must invest in advanced technologies, regularly update and patch their systems, and educate employees about cyber risks and best practices. Additionally, organizations should establish incident response plans to mitigate and respond to cyber incidents effectively.

Furthermore, governments and regulatory bodies are responsible for creating and enforcing cybersecurity policies and regulations. They should promote information sharing and collaboration between public and private sectors to address emerging cyber threats effectively. Governments can also invest in cybersecurity education and training programs to enhance cyber literacy among citizens.

Achieving robust cybersecurity requires a collective effort. Organizations, individuals, and governments must work together to establish a strong cybersecurity culture. By recognizing and fulfilling our shared responsibility, we can create a safer digital environment for everyone.

Why you need a disaster recovery plan

Disasters, both natural and non-natural, can severely impact business operations. As such, businesses must have a well-defined disaster recovery plan in place to mitigate its negative consequences. A disaster recovery plan is a defined set of processes and procedures that outline how an organization will respond and recover from various disasters. It ensures critical business functions can be restored quickly and efficiently, minimizing downtime and reducing financial losses.

The following are some key reasons why businesses need to have a disaster recovery plan in place:

Minimizing Downtime and Loss of Productivity: Disasters can cause significant disruptions to business operations. Without a proper recovery plan, businesses may struggle to get back on track, leading to extended downtime, loss of productivity, and potential revenue loss. A well-prepared disaster recovery plan ensures that necessary measures are in place to minimize downtime, allowing businesses to resume operations as quickly as possible.

Protecting Data and Information: Data is one of the most valuable assets for businesses today. Data can be compromised or lost entirely, leading to severe consequences for a business. A disaster recovery plan should include backup and recovery procedures to safeguard critical data and information. Thus ensuring data can be restored and accessed efficiently to protect the integrity and continuity of business operations.

Ensuring Business Continuity: Disasters can have long-lasting effects on a business if not properly addressed. A disaster recovery plan enables businesses to maintain continuity during and after a disaster. It outlines the crucial steps to ensure essential functions can continue, even in adverse circumstances. By prioritizing business continuity, organizations can minimize the impact of disasters on their operations and maintain the trust and confidence of their customers.

Meeting Regulatory and Compliance Requirements: Many industries have specific regulatory and compliance requirements regarding data protection and business continuity. A robust disaster recovery plan helps businesses meet these requirements and comply with applicable laws and regulations. Businesses can demonstrate their commitment to protecting sensitive information and maintaining operational integrity by having a disaster recovery plan.

A disaster recovery plan is critical to any business's risk management strategy. It provides a roadmap for mitigating the impact of disasters and enables businesses to recover swiftly and efficiently. By investing in a well-designed and regularly tested plan, businesses can protect their operations, data, and reputation, ensuring long-term success in an unpredictable world.

Why security posture is crucial for small and medium businesses

As small and medium businesses continue to grow and thrive in today's digital landscape, it is crucial to prioritize security posture. A strong security posture not only protects sensitive data and information but also safeguards the reputation and longevity of the business.

One of the key reasons why security posture is vital for small and medium businesses is the increasing prevalence of cyber threats. Hackers and malicious actors specifically target smaller businesses, knowing they may have fewer resources dedicated to cybersecurity. Businesses can defend against these threats and mitigate potential damages by establishing a robust security posture.

Maintaining a solid security posture also instills trust among customers and partners. In an era where data breaches and cyber-attacks dominate headlines, consumers are becoming more cautious about sharing their personal information. By demonstrating a commitment to security, businesses can differentiate themselves from competitors and build a reputation as a trustworthy organization.

Additionally, a strong security posture is essential for regulatory compliance. Many industries have specific security requirements that businesses must meet to operate legally. Failing to comply with these regulations can result in severe penalties and legal repercussions. By proactively addressing security measures, small and medium businesses can ensure compliance and avoid costly consequences.

Investing in security posture is a proactive approach that saves businesses from potential financial loss. Recovering from a security breach can be financially devastating, especially for smaller organizations that may not have the resources to bounce back quickly. By implementing robust security practices, businesses can significantly reduce the risk of data breaches, financial fraud, and operational disruptions.

Security posture is paramount for small and medium businesses. By prioritizing security, businesses can protect themselves from cyber threats, build customer trust, comply with regulatory requirements, and safeguard financial stability. Investing in security is not just an expense; it is an investment in the long-term success and resilience of the business.

Is your business a target for Hackers?

In today's cyber threat landscape, no business of any scale is immune to cyberattacks. Hackers constantly scan the Internet for vulnerable targets, and businesses of all sizes can become victims. Here are a few reasons why hackers would target your business:

Valuable Data: The potential for financial gain often drives hackers. Your business becomes an attractive target if you deal with valuable data, such as customer information, payment details, or intellectual property. Hackers can exploit this data for various malicious purposes, including identity theft, financial fraud, or selling it on the dark web.

Industry Reputation: Specific industries are more prone to cyberattacks due to the valuable information they hold. For example, healthcare organizations store sensitive patient data, financial institutions handle large sums of money, and technology companies possess valuable intellectual property. Hackers may target businesses in these industries to gain access to valuable information and exploit their reputation for financial gain.

Weak Security Measures: Hackers often look for the path of least resistance. If your business has weak or outdated security measures, it becomes an easy target. This includes using weak passwords, not regularly updating software, lacking proper encryption, or neglecting employee cybersecurity training. Hackers can exploit these vulnerabilities to gain unauthorized access to your systems and data.

Ransomware Potential: Ransomware attacks have become increasingly prevalent in recent years. Hackers use malicious software to encrypt your business's data and demand a ransom for its release. Any business can become a target of ransomware, especially if they have valuable data and weak security measures in place.

Competitive Advantage: In some cases, hackers may target businesses seeking a competitive advantage. Competitors or individuals with malicious intent may attempt to gain unauthorized access to your business's proprietary information, trade secrets, or upcoming product plans. By doing so, they aim to gain a competitive edge or disrupt your business operations.

Businesses must understand that cyberattack threats are real and can have severe consequences. Implementing robust cybersecurity measures, regularly updating security patches, and educating employees about best practices can significantly reduce the risk of being targeted by hackers. Remember, cybersecurity is an ongoing endeavor. Stay vigilant, stay informed, and protect your business from the ever-evolving cyber threat landscape.

SSL/TLS Certificates and Authorities

As more and more Internet services adopt the SSL/TLS protocol, and Network Box offers various services to secure and protect such traffic; we need clarification and understanding regarding the fundamentals of the protocols - particularly concerning certificates, certificate authorities, and trust. This article aims to clarify this and show how SSL/TLS can be securely implemented. However, note that we will limit our explanation here to server-side certificates and not outbound (client-side) SSL proxying, which is an entirely different topic.

Over the years, the original SSL (Secure Sockets Layer) protocol has morphed into TLS (Transport Layer Security). Earlier versions of the protocol were called SSL; later versions were TLS, but the core ideas and mechanisms behind the protocol are the same. To simplify things, in this article, we'll call it TLS, which fundamentally provides the following:

- Bi-directional encryption of a data stream between a client and server
- Protection against replay attacks, interception, or tampering with that data stream
- The ability for a client to verify that the server is who it says it is
- The optional ability for the server to verify that the client is who it says it is.

The protocol runs on top of another encapsulating protocol - typically TCP/IP (Transmission Control Protocol/Internet Protocol), but may also be UDP/IP (User Datagram Protocol/Internet Protocol) - requiring simply a bi-direction data stream between the client and the server. TLS connections may be directly established (such as to HTTPS port tcp/443) or 'turned on' once an underlying connection has been established (such as the STARTTLS command on an established tcp/25 SMTP connection). Most Internet protocols nowadays provide either a mandatory/alternative TLS port or some mechanism to upgrade a connection to TLS.

Public Key Cryptography

TLS relies on public key encryption. The concept here is that rather than one symmetric encryption key (aka 'password') being used for both encrypting and decrypting, the key is split into two parts - public and private - and it is mathematically infeasible to derive one key from the other. Any of these public/private keys can encrypt data, with the matching pair being used for decryption. The public key can be released or published without concern, while only the private key needs to be protected. In such a system, the public key can encrypt data that only the private key holder can decrypt, or vice-versa. This provides for some interesting approaches - the holder of the private key can prove they have it by demonstrating being able to decrypt data encrypted with the public key or by being able to encrypt data with their private key that is then able to be decrypted with the public key.

This is fundamental. To illustrate, here is an example:

- If User A, the private key holder, wants to prove they hold the key, User B can provide User A with some data.
- User A can encrypt the data with the private key.
- User B can verify that by decrypting it with the public key and comparing the results.

Certificates

At the heart of TLS is the concept of a certificate. Things here can get complex, so the explanation below is simplified and ignores some of the more modern, sophisticated implementations.

Fundamentally, a TLS certificate is a public key, along with some identifying information for the private key holder and a digital signature. This is used for two things: 1) to be able to encrypt data (with the public key) that only the holder of the private key can decrypt, and 2) to be able to verify that the holder is who they say they are. But what is this digital signature? It is produced by digitally signing the certificate data using the private key of some other party trusted by both ends of the communications link. These signatures are protected by the same public key cryptography in that the signer's public key can be used to verify the signature.

An example:

- Certificate A contains the name 'network-box.com' and the public key of services running at Network Box. It is signed by Certificate B.
- Certificate B is a trusted intermediary and contains that intermediary's public key along with its name. It is signed by Certificate C.
- Certificate C is a trusted top-level authority and contains the authority's public key along with its name. It is signed by itself (i.e., the private key of Certificate C).

That example shows just three levels, but there is no specific limit to the number of levels possible. Nowadays, certificates in everyday use have between zero and two or three intermediary certificates before we get to the top-level (sometimes aka 'root') certificate.

Now, say a user wants to make a TLS connection to that Network Box service. The user does a DNS lookup on the name 'network-box.com' and makes a TCP connection to that IP address. The user then goes through a TLS negotiation and typically gets back Certificates A and B. The user then verifies Certificate A (making sure the name in it matches the name 'network-box.com' the user connected to), and ensures it hasn't expired, etc. As A is signed by B, the user can do the same verification on B and do the extra step of verifying B's signature on A by using B's public key. Finally, the user sees that B was signed by C, and the user has a local copy of C (as a 'trusted Certificate Authority' in the user's local storage), so the user can verify C's signature on B using C's public key. This way, the entire certificate chain can be verified based on the trust that both ends place in 'C.'

Trusted Certificate Authorities

In the previous example, C is a trusted Certificate Authority (CA). Typically, these are well-regulated and mutually trusted certificates available to both ends of the connection. The server side trusts C not to sign anyone else as network-box.com without verification, and the client side trusts C to have correctly signed and to vouch for the authenticity of the network-box.com certificate that they signed.

Nowadays, TLS implementations (such as those used in web browsers) include a hundred or more trusted CAs. Each of these CAs have some validation process that they go through to authenticate that someone requesting them to sign a certificate is actually who they claim to be in the certificate. The most common of these validations is domain validation, where they only validate the right to administer that domain, but other more advanced forms of validation are possible (such as company name, etc.).

Typically, in a public network, these root, top-level, trusted certificates are the only ones permitted to be self-signed. All other certificates should be part of a chain of trust leading up to one top-level self-signed trusted certificate. There is nothing stopping you from self-signing a certificate - the issue is getting someone else to trust you.

So now we can see how the TLS protocol works. The client connects to the server, indicates its desire to communicate using TLS, and digitally signed certificates may be exchanged. Typically, a server certificate is always sent to the connecting client for verification, but also, in some cases, client certificates may be sent to the server for mutual verification. The public key cryptography outlined above is used to verify these certificates back to a mutually trusted Certificate Authority.

Certificate Issue/Renew

So now that we've got a good understanding of the system's basic mechanics, let's talk about the creation and renewal of these certificates.

A certificate is created by generating a highly random private+public key pair, putting the requestor's information and public key into a Certificate Request file. That is then sent to the Certificate Authority to issue the certificate. The CA then takes steps to validate the information in the certificate (often just the CN - the Common Name field, usually the DNS name for the service) by validating that the requestor also has control over that DNS domain name. Once validated, the CA re-packages the Certificate Request into a Certificate, signs it with their own private key, and delivers it back to the requestor. As each certificate has an expiry date, the certificate should be renewed before expiry - a process similar to issuing a new certificate except that the original public/private key pair can be re-used.

The requestor will typically validate that they control the domain in the CN of the certificate by one of a selection of validation mechanisms:

- Receiving, and responding to, a secret email to the administrative owners of that domain (proving organizational roles of postmaster, admin, etc.).
- Being able to put a provided secret into the DNS records for the domain (proving administrative control over the domain).
- Being able to put a provided secret into a file in the web server for the domain (proving administrative control over the website for the domain).

It should be noted that certificates can contain just one domain name in the CN, use a wildcard domain (*.network-box.com, for example), or list multiple alternative CNs to cover multiple services in one certificate. Certificate Authorities must validate all such types and all such domains referenced.

Nowadays, there are typically three options for obtaining a signed certificate:

1. Sign it yourself (but only practical in private networks due to the trust issue).
2. Getting a traditional CA to sign your certificate. In such cases, they typically verify just your Common Name, the certificate will expire in a year, and they charge anywhere from US$3 to US$300 per year.
3. Use the ACME (Automated Certificate Management Environment) to issue and renew your certificates. Such certificates are typically issued, renewed, and verified but expire in a month or two, so automatic verification and renewal are essential. A few trusted CAs support this free of charge.

In the April 2023 patch Tuesday, Network Box announced support for the ACME protocol in NBRS-5 - allowing us to support automatically issuing and renewing certificates using this protocol directly. This is particularly important for TLS services offloaded to the Network Box.

Hopefully, by reading this article, one should be able to see how certificates provide the core foundational security of the TLS protocol, how they protect against man-in-the-middle and other such attacks, and also how the trusted CAs become such a security concern (given the vital role they play in domain validation). This month's release of direct support for the ACME protocol in NBRS-5 will go a long way towards simplifying the secure issuing, renewing, and deployment of certificates and implementation of TLS protection.

THE WHITELISTING Approach

There exist two very different and fundamentally mutually exclusive approaches to security:

1. **Blacklisting:** involves specifically blocking what you know to be unwanted and allowing everything else through.
2. **Whitelisting:** involves blocking everything by default and only allowing things you specifically want through.

At the Internet perimeter, we are well accustomed to using the whitelisting approach. Most, if not all, firewall inbound (NET->LAN) policies nowadays block all network ports and only open those ports specifically required for specific permitted services.

But outbound (LAN->NET) at the Internet perimeter, we see more of a mix of approaches. We recommend using the whitelisting approach - block all outbound, and permit only what is explicitly required. But, we still see many customer policies allowing everything outbound, except for a few ports specifically blacklisted.

For inbound email, Network Box has always offered comprehensive policy control and recommended a whitelisting approach - quarantine executables, scripts, etc. (by default), and allow only for specific trusted senders. Most of our customers follow this recommendation with the policies they ask us to implement.

Until recently, using the whitelisting approach on end-points (workstations and servers) has been problematic. By this, we mean blocking any application from executing, except those applications specifically whitelisted. The hundreds of thousands of applications available, each with dozens of interrelated components, combined with frequent updates, often resulted in excessive administrative workload and end-user impact in managing the whitelists. But now, with widespread deployment of code signing on both Microsoft and Apple platforms, combined with more powerful trust rule systems, this is becoming feasible. And it is a practical alternative to traditional blacklist-based approaches such as host-based anti-virus and intrusion prevention systems.

Modern host-based whitelisting systems are flexible (cloud-based, with powerful rules supporting application signatures, digital signing certificates, as well as metadata such as file paths, parent application, etc.) and easy to deploy with inherited trust mechanisms. They finally offer a viable alternative to traditional host-based anti-virus systems. They also go beyond merely stopping the latest ransomware attack, to cataloging and reporting on applications actually running on the hosts in your network - providing for effective per-host and per-user policy control. They are still more complex to maintain than traditional anti-virus systems, but combined with a managed service, are finally becoming something truly useful and perhaps the ultimate solution to secure end-point devices.

Whitelisting Signatures

At the core of a whitelisting product is the endpoint engine. This obtains signatures of objects being executed, compares them to the database of signatures listing what is permitted (aka 'the whitelist'), then enforces and produces audit logs. But what is a signature?

Most commonly, one-way hash functions have been used. These are a form of one-way encryption, taking a large object, then applying a mathematical algorithm to reduce it to a much smaller 'hash' value. For example, the MD5 hashing algorithm takes any arbitrarily large object and produces a hash of just 128 bits. You cannot take a hash and reconstruct the original object (hence the term 'one way'), but you can simply compare the hashes of two objects and, if identical, deduce that the objects are the same. The hash collision rate (where two objects produce the same hash value) must be incredibly low for this to work. For security purposes, it should also be extremely hard to force a collision (by adjusting an object to make it produce a known hash).

The whitelisting approach is thus to take fingerprints of all permitted objects and store them in a list. Then, whenever an object is to be executed, we can compare its hash against our whitelist and permit/deny it as appropriate. Such an approach is very secure but has one critical vulnerability - if the attacker can adjust his malicious code to have the same hash value as a whitelisted (presumably common) application, then it will be permitted to be executed. While computationally hard to do, this is not impossible, and in recent years more and more hashing algorithms have fallen vulnerable to such approaches.

The approach Network Box has chosen is to take multiple fingerprints, using a selection of five of the most secure hashing algorithms, to produce a 'handprint' for each executable object. The chances of one of these algorithms being compromised are small, but the complexity of all five being compromised is so infinitesimally tiny as to be practically impossible.

Note that signature technology can typically be applied to the objects or the certificates used to sign those objects (in the modern world of code signing certificates) - to minimize issues with installing application updates. But bear in mind the additional risks of trusting a particular developer entirely).

How to Deploy Whitelisting?

So how to deploy whitelisting on a workstation/server?
There are three common approaches:

1. Learning Mode: With this approach, all the existing executables on the machine are pre-indexed and trusted, and the machine placed in learning mode will automatically trust new executables run during the learning period. The advantage here is simplicity - the whitelist is built automatically and is pretty complete. The disadvantage is obvious - not all pre-existing executables on the machine may be desirable (either from the point of view of policy control or simply because they may be malware).

2. Monitor Mode: Here we monitor all executables run on the machine over time, but instead of blocking, we merely alert and have an administrator review, categorize, and decide to permit/deny in the whitelisting policy. The advantage here is that the resulting whitelisting policy is very complete, but it does require administrative effort.

3. Pre-Trusted: The pre-trusted approach builds upon pre-existing lists of known common popular applications to be permitted by default, with everything else denied. The advantage is that it is quick and simple to implement, but the disadvantage is that anything custom or unusual would not typically be trusted by default.

Which is the best approach?

In Network Box's view, the Learning Mode is simply not a good solution as it runs the risk of permitting malware into the organization during the learning period. Depending on the situation, we typically recommend a combination of **Monitor Mode** and **Pre-Trusted**. This balances the benefits of minimizing the deployment period while maximizing security.

Is Whitelisting for Everyone?

The whitelisting approach is not a universal solution suitable for everyone. Like most things involving IT security, there is a trade-off of security vs. convenience, and the whitelisting / blacklisting approach seems to exemplify that.

IT Administrators are very familiar with traditional blacklisting anti-malware systems. They deploy them onto the network, and within minutes, pre-existing malware is identified and ongoing protection provided. However, such systems are proving to be increasingly hard to maintain their effectiveness. The explosion in the sheer volume of malware, and the increase in ransomware attacks for commercial gain, have seen more and more malware get through to infect systems. It doesn't matter how good an anti-malware blacklisting system is and how much malware it blocks; it only has to let one single thing through to result in a nightmare.

Whitelisting provides an alternative to this, which is probably as close to 100% secure as can be realistically achieved. Whatever the malware authors come up with is blocked by default, but that comes with the downside that whatever new applications or updates you require on your network will also be blocked by default. Whitelisting introduces an authorization step into the flow of application installation/update, which can be good or bad, depending on your viewpoint.

That said, one very positive side-effect of deploying whitelisting is the increased visibility of what applications are being run on your workstations/servers and by whom; that comes from the audit logging and policy control. Whitelisting can be used to block not just unwanted malware but also non-core applications from your network.

The Network Box Whitelisting Solution

Network Box has partnered with Whitecloud Security to bring managed zero-trust endpoint security to our customers. With the mature and well-established Whitecloud Security endpoint, using unique five signature handprint technology, and combined with trust architecture and Network Box Managed Security Services taking on the administrative overhead.

The approach involves a small zero-trust endpoint agent installed on Windows servers/workstations, connecting to a cloud-based service for management and control. The usual deployment approach is:

1. Identified suitable endpoints to be protected. This can be all endpoints or only those at high risk (such as accounts/finance workstations, laptops leaving the office, servers, etc). These endpoints are grouped so that policies can be applied to the group (whitelisting an application on one workstation automatically whitelists it on all others in the group) or for specific workstations.
2. Deploy zero-trust endpoint protection in Monitor mode. During the monitoring period, Network Box SOC engineers will adjust trust policies as necessary and alert as to any malicious activity detected.
3. After a suitable monitoring period (usually one to two weeks), enable enforcing of the policy.
4. Network Box SOC engineers provide ongoing monitoring and support to enforce the whitelisting policy.
5. Customer Administrative staff have complete visibility and control over the above process, working alongside Network Box SOC engineers.

This technology also brings with it an interesting solution to pre-existing infected machines (ransomware, etc). In such cases, the endpoint agent can be installed on the infected machine in Pre-Trusted mode (only trusting base Microsoft and other common applications), and the machine then brought safely back online for investigation and recovery (safe in the knowledge that nothing untrusted will be permitted to run).

With a simple per-device per-month cloud-based pricing model, bringing zero trust endpoint protection to the Network Box customer base provides us exciting opportunities to help customers with their high-risk endpoints, serviced from our existing Managed Security Services frameworks.

Managed Zero-Trust End-Point Security

With the Network Box managed Zero-Trust End-Point Security solution now available and released globally, we thought showing some example deployment case studies might be helpful. In this article, we present three case studies illustrating different deployment approaches for whitelisting technology.

Case 1 - The Reactive

The user here has a network of approximately fifty workstations, laptops, and servers. All run traditional antivirus technology. A salesman took his laptop out of the office and got infected with some trojan malware while on a business trip. Upon returning to the office, a ransomware application was downloaded and executed by remote hackers - encrypting his laptop and files on several network shares. The network admins have already disconnected and isolated what they could but are concerned that with the trojan application in place - hackers have remote access to the network for lateral spread. Taking everything offline, forensic imaging, and one-by-one cleaning things up would take an estimated 7 to 10 days, with associated business impact costs.

In cooperation with Network Box SOC, the following actions are taken:

- All suspected infected machines are taken offline and rebooted into safe mode. Zero-trust End-point security is installed from USB, and machines rebooted into a group-based application control policy only permitting a very limited set of pre-trusted applications to be run (primarily Microsoft and some business-critical applications). At this point, these machines are safely brought back online, used as normal, and critical data is extracted.
- A Network Box device is placed at the Internet perimeter to replace the simple firewall there before (with zero outbound policy), and effective policy rules are put in place to control both inbound and outbound traffic. Infected LAN, IDS, and IPS engines are enabled to monitor and control outbound traffic.

With the hackers locked out of the network and the infected machines back under control, the business impact is limited, and ransomware-encrypted files can be restored from backups.

Case 2 - The Cautious

Here, we have a large network of several hundred workstations, laptops, and servers. Traditional antivirus technology is run on these machines, and the network is protected at the perimeter by a Network Box. The owners and administrators are concerned that an end-user will make a mistake and click on something they shouldn't - potentially bringing down the entire network.

We identify key high-risk workstations and servers, including:

- Internet-accessible servers running web, email, and collaboration software
- Accounts workstations
- Key decision-maker workstations (including high-level executives, the financial controller, etc.)
- Out-of-the-office laptops

Zero-Trust End-Point Security is deployed to all those high-risk machines and runs in monitoring mode for two weeks. During that time, Network Box SOC staff monitor the applications being run and whitelist as necessary. Some potentially unwanted applications are identified, and Network Box SOC staff work with the admins to address these on a case-by-case basis. Towards the end of the two weeks, the number of alerts raised for unrecognized applications falls to zero, and the machines moved to enforcing mode (blocking the execution of untrusted applications).

The approach here is not perfect, and particular care needs to be taken regarding network shares accessible by end-points not protected by zero-trust (as ransomware infections on those end-points could encrypt files on the network shares). Security can never be 100%, and there is always a balance between convenience, cost, and security; such a risk-based approach attempts to address that balance trade-off.

Case 3 - The Prepper

This is a relatively small network. A financial services firm with a small number of highly paid staff offering consultancy services. Key decision-makers are concerned that a ransomware attack, or network intrusion, could leak sensitive customer data (particularly given that most staff use laptops that spend time outside the office network protection).

Zero-Trust End-Point Security is deployed to all workstations, laptops, and the network server, in monitoring mode. Over a period of two to three weeks, Network Box SOC staff monitor the applications being run, whitelisting as necessary, until the machines can be moved to enforcing mode (blocking the execution of all untrusted applications).

During the deployment and subsequent months, several unwanted and potentially dangerous applications are blocked from running on the network. Network Box SOC staff alert the office manager for follow-up with the end user. The key decision-makers are impressed with the reports they can obtain showing which applications are being run by which users at what times.

Conclusions

Moving from a blacklisting (antivirus) to whitelisting (zero-trust) approach requires a shift in mindset. Each approach has its advantages and disadvantages, best summarized as:

Whitelisting vs Blacklisting Pros and Cons

Metric	Whitelisting	Blacklisting
Effectiveness against known malware	100%	Close to 100%
Effectiveness against emerging malware	100%	Perhaps 90% to 95%
False positives	Updates, and new installs	Few
Maintenance of the list	Admin or SOC managed	Vendor
Blocking action	On execution	On download / scan
Visibility of applications used	Full reporting	Typically none
Policy control of applications used	Full control	Typically none

You can see that the biggest drawback of the whitelisting approach is that it requires end-user / admin maintenance of the whitelist. At the same time, the most significant differentiator (apart from anti-malware effectiveness) is the improved reporting and control over which applications are permitted to be run on the network. By simply not trusting (or adding to the whitelist) unauthorized (as opposed to malicious) applications, effective policy control can be implemented. Whitelisting gives the end-user full control and reporting on which applications are actually being run on their end-point devices.

The Network Box approach solves the end-user administrative burden problem of maintaining the whitelist by moving that function to Network Box SOC engineers and our managed service. We offer self-managed, SOC-managed, as well as hybrid combinations. In other words, we offer all of the advantages of zero-trust with none of the drawbacks.

Issues with DNS

(Domain Naming Sytem)

If your network is like most companies, you are likely using Active Directory and therefore have Domain Controllers (DCs). Your workstations are likely using these DCs as their DNS servers. The DCs, in turn, have a configuration for DNS forwarders, which are used to resolve public IP addresses. In the majority of cases with our clients, these forwarders are configured to be the DNS servers of their primary ISP. This configuration might have been viable in the past, but today it is not advisable. In fact, it is discouraged. And here is why.

The ISP DNS servers are not a public service

They do not respond to DNS queries from the Internet. They respond to your servers because you are using their network. In other words, your public IP is authorized to query those servers. So, what happens when you change ISP? Suddenly, you can't get to the Internet, and you scratch your head thinking, "the new ISP did something wrong," when instead, you have a simple DNS issue - your new public IP address is not authorized to query those servers.

Of course, the first solution you'll think of is to replace those forwarders with the new DNS IP addresses the new ISP provides. But what happens if you have 2 ISPs, in load balance or high availability? Do you change the forwarders every time you fail over to the 'other one'? Or do you configure the DNS servers recommended by both ISPs? If you do that, when you're using the 'secondary' DNS server, you'll see delays, and your users will complain of slow Internet, which is very close to saying that the Internet isn't working.

Consider that a DNS query can take hundreds of milliseconds, and the secondary DNS server is only queried once the primary times out. Then consider how many domain resolutions your browser needs to perform for each web page you're visiting; you'll quickly realize that DNS malfunction is a very likely cause of Internet sluggishness. In my direct experience, DNS issues are by far the most frequent reason why the Internet is slow.

DNS can be attacked

Another very important reason to avoid adopting such forwarders is that DNS can be attacked. Several attacks can be carried against DNS servers, but here I am referring more to DNS spoofing, whereby your workstation ends up requesting DNS from servers that aren't the ones you think you configured. This is a catastrophic attack because now, every query your browser runs gets a reply that will point your browser to the IP the hacker wants you to reach; the results of that can be catastrophic. Protecting your DNS with a DNS proxy is essential because it avoids this attack altogether.

Rate Limiting for Queries

Many bypass the ISP DNS issue by using public DNSs. I see many using 8.8.8.8. While there is nothing wrong in doing so, those IP addresses belong to Google. So the first thing that you need to consider is that you're telling Google every domain you're visiting. Is that something you want? Or did you think that Google has a big heart and is providing you with that service for free? Aside from this 'conspiracy theory,' this IP address (and its companion 8.8.4.4) apply rate limits to how many queries per second or minute they will allow. If you're a small organization, this may not really matter, but it is quite easy to reach that limit - after which they stop responding for a period of time. And again, you'll think there is something wrong with the Internet when it's just your DNS configuration causing you problems instead.

Cloudflare also offers a similar service with their 1.1.1.1 IP address. Many years ago, the DNS of choice was 4.4.4.4 and 4.4.2.2. But then Layer 3 bought the entire 4.4.0.0/16 subnet. The DNS servers are still working, but Layer 3 is not in the business of providing free public services. Before you know it, they might well take those servers down. And again, you'd be stuck with no Internet.

So what is the solution to all these possible issues?

If you are a Network Box client, the solution is to point your forwarders to the local Network Box LAN IP (or DMZ or whatever). We will configure all our devices to run their own secure DNS service (DNS server 127.0.0.1) and use the Internet ROOT DNS servers as the primary forwarders.

There are many reasons why this is a good idea. The ROOT servers are those operated by the registrars, so when you make DNS configuration changes, they are the first to know, and the propagation has to start from here. I have seen cases where 8.8.8.8 took hours to note a DNS change when the root servers had already seen it within minutes of the change being made. The root servers are by far the most reliable. They are run by the registrars and 'know,' which are the authoritative servers for a domain. At the end of the day, every other DNS server ends up querying these servers first to find the IP of the DNS server that is authoritative for that domain. So, why not query them directly? And by using the Network Box LAN IP as your DNS forwarder, you will also avoid the DNS spoofing attack, which is a big thing to consider.

If you are a Network Box client and you are also using a DNS proxy service such as OpenDNS, you may want to reconsider how you're spending your budget. The service they offer is something you already have with Network Box. Why spend the money twice?

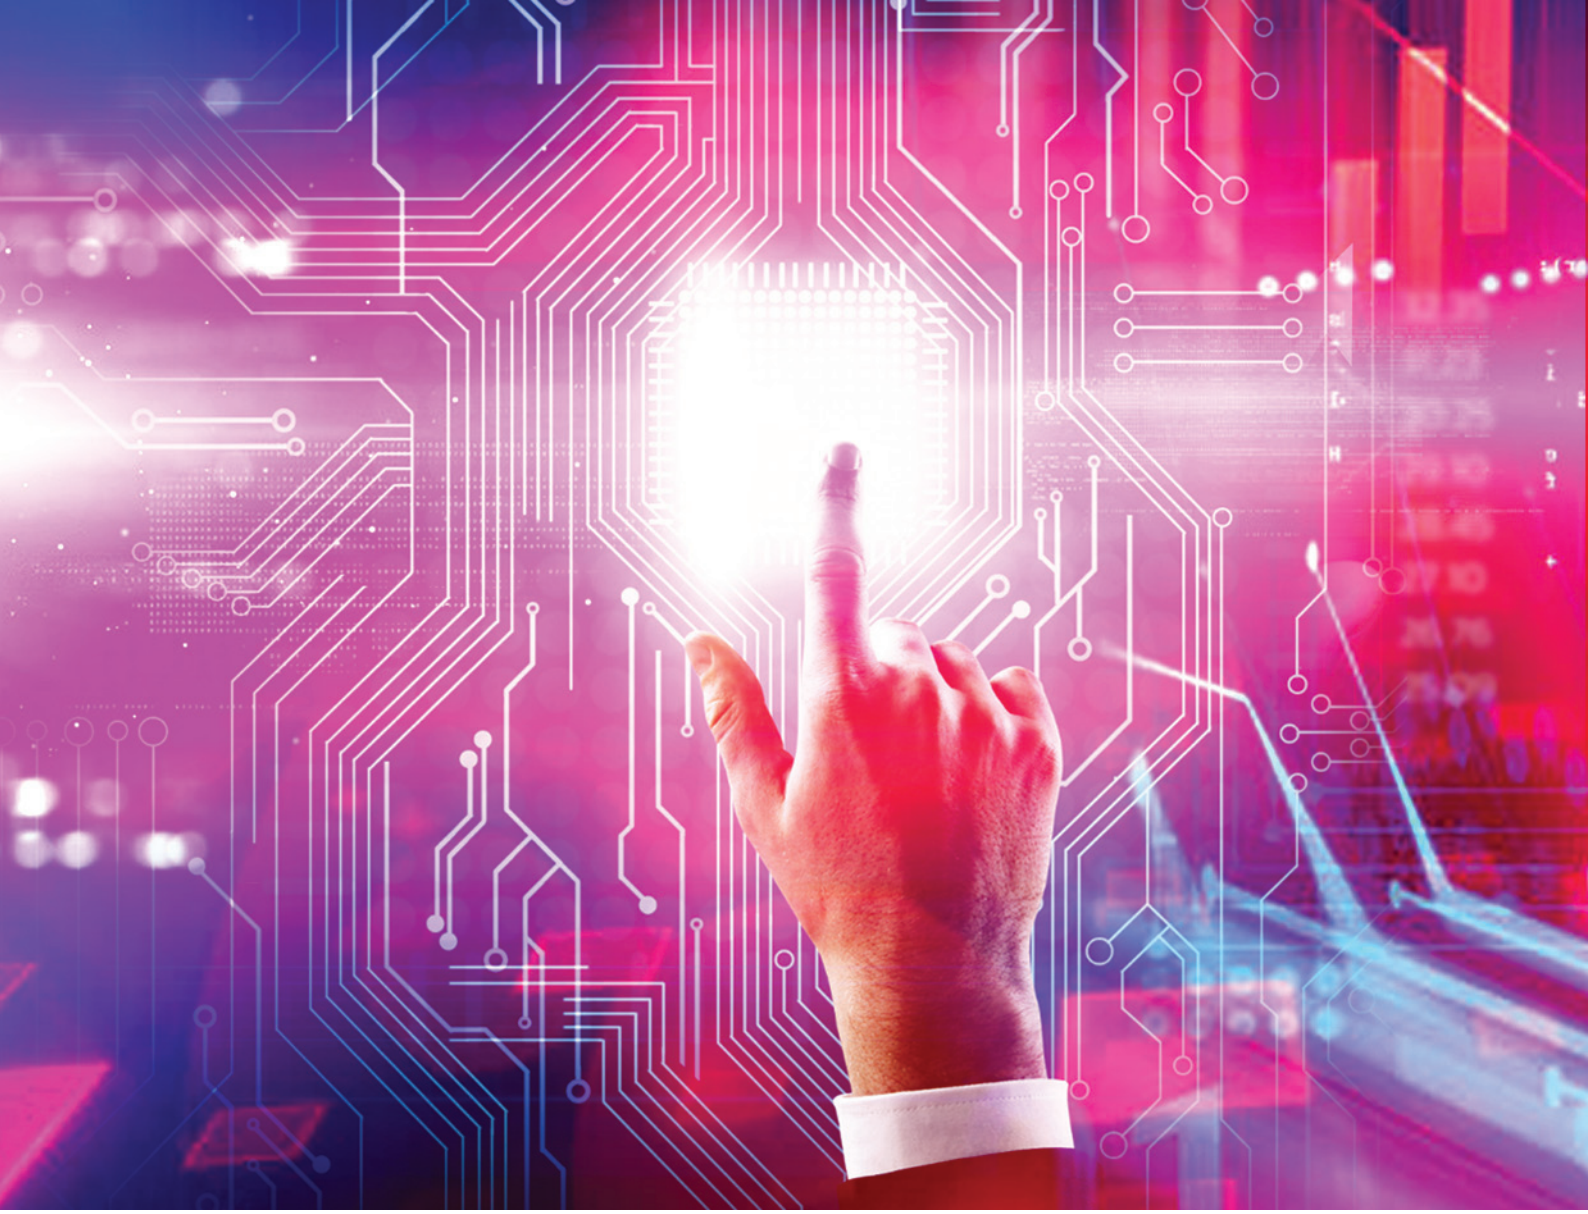

Configuration Reviews

Most security frameworks include periodic configuration reviews as a core requirement. Whilst all configurations should adhere to the defined security policy at initial deployment, and all subsequent changes should have been made in accordance with that policy, this is often insufficient.

For example:

1. Individual configuration changes may impact other configuration items in unexpected ways (such as a network addressing/routing change exposing firewall rules to new traffic).
2. Policies and risk tolerance may change so that what was acceptable in the past may no longer be acceptable today (such as new threats, vulnerabilities, and attack vectors).
3. Staff may leave, and with them, the knowledge of mitigations previously put in place (such as the reason for a particular service to be exposed and steps taken to mitigate that risk).

Network Box has always followed the approach that the customer sets the policy, and the SOC securely implements that policy in the configuration. We are frequently asked to recommend policies or to suggest optimal deployment approaches, but the policy itself is entirely under the control of our customers.

Previously, we formalized our general security recommendations by introducing a set of best practices (**https://network-box.com/best-practices**); developed over two decades of delivering Managed Security Services, investigating security incidents, and working with our customers to protect their networks. These Best Practices represent the most common forms of network infiltration and data breaches that we see affecting networks worldwide. Many of these best practices can also be found in common standardized security frameworks. Our Security Engineers refer to these Best Practices when designing defense systems for networks under management when processing policy change requests, and during periodic configuration reviews. While ultimately, the customer decides the policy; we strive to inform, warn, and point out when policies conflict and open up networks to common attack vectors and unnecessary risk.

As part of the general move towards Managed Detection and Response, Network Box SOCs have begun conducting ***formal configuration reviews*** with reference to these best practices. You will see the results of these reviews as PDF reports attached to tickets raised in Box Office / NBSIEM+. The initial review will highlight all areas of concern, and subsequent reviews will also include a table of changes (additions, changes, and resolved concerns) since the previous review.

We encourage you to work with our Security Operation Centres to address highlighted items and to use this system to improve your security policy and defense.

BARRACUDA
ESG Zero-Day Vulnerability

In mid-May 2023, Barracuda (a manufacturer of network security appliances) discovered unusual traffic coming from some of their ESG (Email Security Gateway) appliances. These appliances filter email for viruses/spam and can be deployed as physical or virtual machines. Barracuda followed this up on 30th May 2023 with public disclosure of the issue labelled CVE-2023-2868 - a remote command injection vulnerability with evidence of in-the-wild exploitation, back to at least October 2022.

In their disclosure announcement, Barracuda revealed that they had already released patches on 20th May, which initially seemed more of the same (just another vulnerability, another exploit, and patches to address the issue). However, on 6th June, Barracuda shocked the security industry with an update saying that all impacted devices should be completely replaced (not just patched), irrespective of firmware or patch level. Such a global recall was unprecedented and indicated a problem far more severe and deeply embedded than first thought.

CVE-2023-2868

The CVE itself sounds fairly malicious:

> **CVE-2023-2868**
> A remote command injection vulnerability exists in the Barracuda Email Security Gateway (appliance form factor only) product affecting versions 5.1.3.001-9.2.0.006. The vulnerability arises out of a failure to comprehensively sanitize the processing of .tar file (tape archives). The vulnerability stems from incomplete input validation of a user-supplied .tar file as it pertains to the names of the files contained within the archive. As a consequence, a remote attacker can specifically format these file names in a particular manner that will result in remotely executing a system command through Perl's qx operator with the privileges of the Email Security Gateway product. This issue was fixed as part of BNSF-36456 patch. This patch was automatically applied to all customer appliances.

Mandiant has penned a thorough analysis of the issue for those interested in the more technical aspects:

https://www.mandiant.com/resources/blog/barracuda-esg-exploited-globally

In brief, when the affected Barracuda appliances receive an email containing an attached 'tar' (Unix/Linux Tape Archive) file, it attempts to extract the contents for further analysis. A flaw in the Barracuda code passes the list of filenames unsanitized as arguments to a system command, giving the attacker control over the command actually executed by manipulating the filenames of files in the archive.

Why so serious?

Exploit of this vulnerability provided attackers with complete control over the affected appliance. As such appliances often contain credentials for access to other network equipment (such as LDAP, FTP, and SMB servers), the attacker can exploit other machines on connected networks using remote access. With full access to the Barracuda appliance, attackers can also install backdoors, proxy tunnels, and a kernel rootkit to compromise the appliance completely.

Given the level of compromise, Barracuda had no choice but to recommend a complete replacement of affected appliances. They simply could not be sure that a simple patch could remove all remnants of all exploits.

While including such a fundamental weakness in a shipping security appliance was undoubtedly careless, Barracuda can be applauded for handling the follow-up in an open and responsive manner.

It is hoped that this event becomes a wake-up call to everyone in the network security community. While we are used to seeing the vulnerability-exploit-patch cycle, we must be aware of other consequences of exploits and how bad they can be.

STRENGTHENING CYBERSECURITY:

Why Government Legislation is IMPERATIVE

The need for robust cybersecurity measures cannot be overstated in today's hyper-connected world, where technology has become an integral part of our everyday lives. From our omnipresent mobile phones to our laptops and desktops, to smart devices such as CCTVs, refrigerators, and webcam-equipped televisions, which rule our day-to-day existence - everything is an internet-connected computer now. With cyber threats constantly evolving, posing significant risks to individuals, businesses, and even national security, it is critical for governments to enact legislation to tackle these issues head-on.

Given the objective failure of organizations to secure themselves from hackers and malware, government legislation on cybersecurity is necessary, bringing potential benefits to society as a whole. Just look at the number of confidential credentials posted on the Dark Web by hackers, which stands at 12.6 billion and counting. There are more hacked accounts than there are people on Earth. **If that is not a call to action, I don't know what is.**

Safeguarding personal information

In this digital age, personal data is constantly at risk of being compromised. Yet governments and organizations force us to give up more and more of our information. We often have no choice but to fill in the online forms presented to us, typically with the exact information a hacker can use to steal our identities. Identity theft, financial fraud, and unauthorized access to private information have become alarmingly common. Government legislation on cybersecurity can empower individuals by instituting standards and regulations to protect personal information. Implementing robust data protection laws, such as stringent encryption protocols and mandatory breach notification requirements, can significantly reduce the risk of data breaches and protect citizens from the potential consequences of cybercrime.

Educating and enhancing public awareness

With the rapid advancement of technology, cyber threats are continuously evolving, necessitating ongoing education and awareness initiatives. Government legislation in cybersecurity can facilitate the implementation of public awareness campaigns, educational programs, and training opportunities to increase citizens' cyber literacy. Helping citizens become aware of the tactics used by cybercriminals is imperative. By promoting responsible digital practices and equipping individuals with the skills to protect themselves online, government legislation can empower citizens to navigate the cyberspace securely, ultimately reducing susceptibilities to cyberattacks. Artificial intelligence is also bringing a whole new level of threat, as what we see/hear/believe is being challenged with ever more sophisticated deep fakes.

Supporting economic stability

Cyber threats not only jeopardize individuals' privacy but also pose a significant risk to our economies. Businesses of all sizes, from multinational corporations to small startups, are increasingly vulnerable to cyberattacks that can result in financial losses, reputational damage, and even bankruptcy. Government legislation, in the realm of cybersecurity, can foster a secure environment for businesses to thrive. Governments can provide businesses with the necessary tools to safeguard their digital assets and ensure economic stability by mandating adequate cybersecurity measures and promoting information sharing about emerging threats. For governments to implement threat intelligence and install 'cyber radar' to monitor threats in real-time, would make all the difference to ongoing economic stability.

Protecting national security

Cyberattacks now have the potential to disrupt essential services, compromise sensitive government information, and even threaten national security. By legislating cybersecurity, governments can establish comprehensive frameworks to protect critical infrastructure, safeguard classified data, and respond effectively to cyber threats that may originate from internal and external sources. This proactive approach allows governments to counteract potential attacks and reduce the impact on the nation's security. The first blow to a nation's security, even in the case of a war commencing, is far more likely to come from a targeted cyberattack than a barrage of cruise missiles. Indeed, modern warfare now includes the use of hackers and malware, as much as tanks and aircraft. The biggest threat to a nation or an economy is likely the use of an enemy's cybersecurity equipment during a time of peace, only for that equipment to become a Trojan Horse if and when a war, or even a cold war, commences.

Promoting international cooperation

Cyber threats are not confined within national borders; they are a global concern. Government legislation on cybersecurity creates a foundation for international cooperation in combating cybercrime. On a non-military, law enforcement level, global collaboration can help the entire world combat cyber criminals much more effectively. By establishing international standards and frameworks, governments can collaborate with other nations to address cross-border cyber threats more effectively. This approach will facilitate information sharing, joint investigations, and the extradition of cybercriminals, ultimately leading to a safer and more secure cyberspace on a global scale. In the end, there is only one Internet to police, despite that Internet existing across some 206 economies. This means securing the Internet needs to be done collectively. It is simply impossible for one country or economy to do it all alone.

The urgency to prioritize cybersecurity has never been greater, with cyber threats escalating in complexity and severity. New malware, vulnerabilities, and hackers appear all the time. They target our identities, our assets, and even our core beliefs. Unfettered attacks on societies can, and unfortunately do, result in a world where not even what is fact and what is fiction is clear anymore. Facts matter. Truth matters. The government's role in legislating cybersecurity cannot be underestimated. Leaving all of this to companies, organizations, and private individuals just doesn't work.

By enacting comprehensive cybersecurity legislation, governments can protect national security, safeguard personal information, support economic stability, promote international cooperation, and educate the public about the importance of cyber resilience. Through these measures, governments can create a safer and more secure digital environment for individuals, businesses, and nations. The time to act is now, and through collaborative efforts between governments, industries, and citizens, we can build a resilient cyber infrastructure that protects us, empowers us, and propels us forward into a secure digital future.

Scanning and the External Threat View

When analyzing the security posture of a computer network, various viewpoints can be considered, of which the top three are usually:

1. The Internal View
Which services are reachable for a potentially malicious intruder accessing them from within the LAN/DMZ.

2. The Privileged External View
Looking at services reachable to external privileged partners on the Internet - usually accessing from specifically privileged source addresses, MPLS networks, or via VPNs.

3. The Public External View
Services reachable to the general Internet.

While the public external view is not the only concern, it is commonly the most likely vector for a breach/intrusion. Thus, it is a focus for many protection technologies and policies.

To understand your security posture, it is crucial to have a clear view of what hosts and services are exposed to each of those viewpoints. While configuration reviews can go some way towards helping network reconnaissance by scanning (including enumeration of reachable hosts and services, and attempted identification of these), it is still the most effective technique.

Network Box External Threat View

Network Box Security Response has launched the **Scan External View** cloud service that operates as follows:

- Firstly, we need to know what to scan. To do this, we build a list of public and private IP addresses, domain names, and other such information for each asset under management. These 'asset attributes' are maintained automatically by parsing Network Box configurations but can also be manually administered (for attributes not directly visible in configurations). Administrators and SOC engineers can view these attributes on the Asset screen of NBSIEM+.

- Periodically (once a week or after major configuration changes by default), we comprehensively scan UDP and TCP ports on all public IP addresses from sources on the public Internet. This scan is typically in four parts:

  1. Scanning:
  for open UDP or TCP ports and retrieving welcome banner messages from these reachable services.

  2. Service Identification:
  based on banner analysis and other fingerprinting technologies.

  3. HTTP/HTTPS Identification:
  specifically looking for web services.

  4. Basic Common Vulnerability Identification:
  highlighting Best Practices findings.

- The results of the scan (discovered hosts, services, and vulnerabilities) are stored in a database and made available in the NBSIEM+ **Asset > Scans** screen, as well as for reporting purposes.

This is not intended as a full Vulnerability Scan. It is purely a reconnaissance scan, only showing what services (protocols/ports) and hosts (IP addresses) are open and visible to the public Internet. The scan is lightweight and only issues requests commonly seen daily in such Internet traffic.

Usage of the Results

The results are primarily used by Network Box SOC engineers as part of the configuration review process. They are part of a consistency check to ensure that the configuration correctly reflects the customer policy.

Network Box Security Response engineers also use the database when handling emerging vulnerabilities. We can quickly search for affected services and identify networks under management with those services reachable from the public Internet.

The Network Box *Scan External View* cloud service has been released and is now in operation globally. The results of these scans will be available to customers later this summer in the next release of NBSIEM+. This is the first of several *Network Box Red Team* services to be offered.

Artificial Intelligence and Machine Learning

In recent years we have seen the gradual introduction of Artificial Intelligence and Machine Learning (AI/ML) technologies into our everyday lives. From talking to our Siri/Alexa/Google Home devices, to automated chat response systems, computer vision, and self-driving cars - these new systems are no longer 'programmed' procedurally. Instead, they are 'taught' or 'trained' in what is expected and respond with 'how' to do it, decided by the machine model itself.

We've grown accustomed to the predictability of computerized systems - given the same input, the same outputs will be derived time and time again. 2+2 will always equal 4. But these new AI/ML systems behave much more randomly - providing the ability to adapt to changing inputs - sometimes impressing with their comprehension of what we are asking, but also dramatically failing in bizarre ways.

As with all such tools, the technology has both good and bad sides. In this month's article, we'll talk about the positives of AI/ML by providing three examples of how it is being used today for Computer Security.

1. Access Denied Security Events Analysis

For decades, we've been using Heuristics to analyze access denied security events. An example would be setting a threshold for network port access denies per minute and alerting/blocking should that threshold be exceeded. The classic 'Portscan' deny.

The problem with this approach is twofold:

1. The threshold must be manually set and tuned depending on the individual network configuration
2. Slow scans (where the attacker deliberately scans very slowly) are not detected.

These types of heuristics are classic examples of procedural programming - **if this, then that.**

AI/ML models provide an alternative approach. Here, we train the model with examples of normal access denied traffic and targeted attack traffic. We teach the model by example and have it set the thresholds automatically based on that training. Like a child, the computer learns - we don't tell it how to detect a targeted attack, but merely train it to what such an attack might look like. After training, we can then feed a stream of real network events into the model, and it can tell us if it sees anything that looks like an attack worth responding to (so that we can alter/block/respond appropriately).

This approach can be used not just for port scan detection but also for more general high-level access denies such as application logins, detecting brute force, or user enumeration type attacks.

2. General Behaviour Analysis

While heuristics have worked well for access denied security events, they haven't been generally useful for network behavior analysis. The idea here is to set thresholds and criteria for what normal network traffic might look like, so we can alert on anything abnormal. There has been some success here with protocol enforcement (such as defining what particular packet types for a specific protocol might reasonably look like), but such a whitelisting approach is laborious and must be customized for each and every protocol and application.

AI/ML holds great promise for this. Rather than procedurally programming the behavior and thresholds for each and every protocol, we merely train the model with known good behavior and have it alert on anything different.

3. Meta Analysis

While general behavior analysis looks at protocols and applications, meta analysis looks at network traffic attributes (such as the source and destination IP addresses, authenticated users, countries, networks, times of day, etc.). Here, AI/ML can be trained with normal network traffic and alert on anything different. An example of this would be network logins on a Sunday from users who typically work Monday to Friday.

Despite the meteoric rise of ChatGPT, AI/ML is still in its infancy, particularly with respect to its use in computer security. Computers have historically been most useful in situations with clearly defined inputs, outputs, and procedural processes - and have struggled with more vague problems such as pattern matching. AI/ML is more 'fuzzy' and the requirements less well defined - the main issue being false positives. AI/ML often impresses with its accuracy but equally often fails dramatically for no discernable reason.

Network Box Security Response continues to work deploying AI/ML models at the moment, primarily to our NBSIEM+ Event Analysis and Incident Response systems. Over time, we expect this tool to become more useful for this and start to be deployed to perimeter gateway protection and endpoints.

NETWORK BOX

TECH REVIEW 2023

Contents

www.network-box.com