THE MICHIGAN
MOONSHOT

LET'S EXPAND COMMUNITY NETWORKS IN MICHIGAN

A Michigan Moonshot Whitepaper

Securing Community Networks:

20 CIS Critical Security Controls

Table of Contents

Introduction

As a community network operator, you are no stranger to how important broadband and network access is to your communities and constituents. You help underserved residents and help propel them through the 21st century, enabling them to access modern services and be competitive in the worldwide marketplace. Your networks and the services you provide are important, trusted, and required for us to thrive.

Unfortunately, being connected to a worldwide audience is a two-way street; it also opens up the potential and opportunities for attackers located worldwide to take advantage of the systems of not only your organization, but your customers as well. Threat actors around the globe can be a gigantic risk to your network, whether they are unskilled and taking advantage of simple opportunities, or whether they are advanced nation-state attackers with specific political and monetary goals. Even if these malicious goals focus on one of your customers, you may find yourself in their crosshairs as the "weak link" in one of your subscriber's cybersecurity posture.

Because of their relatively smaller size, community networks typically do not have the same level of cybersecurity defenses than those available at larger or for-profit organizations. Attackers the world over do not care that we are smaller or lesser resourced; we are all equal potential victims in their eyes and we have to be able to react and respond just as well as any other network operator out there. Complicating this is the view that cybersecurity can be extraordinarily complex, as there is the potential for attackers to exploit any component under our control. This includes our desktops, servers, WAN equipment, integrations, cloud providers, and our employees themselves. Properly understanding how all the components of your network interact can be daunting.

Thankfully, solutions are here and can be surprisingly simple. This security framework is based on the Center for Internet Security's "Top 20 Security Controls" and tailored towards the unique elements that need to be addressed as community network operators. By following this framework, you can elevate your cybersecurity program to a robust and extremely defensible position, able to resist even the most complex and involved attacks. Each topic will be introduced to you in how it is important to your operations, along with an honest look on how an attacker would take advantage of you in that particular situation. Simple solutions are provided to address each of these concerns, and where possible low-cost or free options and processes are included to help jump-start your security initiatives.

As we move forward with keeping our networks and customers secure, one vital point to make is that **together, we win**. Communities such as these are critical in sharing information, ideas, concepts, and technologies so that we can act as a unified force against the wide assortment of attackers wishing to do us harm. Remain active in talking with your peers and sharing ideas, as attackers' biggest weakness is understanding their attacks so that they can be defended against and vanquished.

Kevin Hayes
Chief Information Security Officer
Merit Network, Inc.

Inventory and Control of Hardware Assets

To create and implement your cybersecurity program, the most critical aspect is to fully understand the hardware assets you are trying to protect. All of our efforts to protect our environment are wasted if there are servers, laptops, switches, routers or wireless access points that we simply do not know about and do not know have to be secured. As a network provider, your hardware is your lifeblood, and ensuring that you have a full and complete understanding of your assets is a requirement for not only providing services to your customers but also making sure you can provide them in a secure and reliable manner.

THREATS

Being able to directly attach to your network is an attacker's best-case scenario, as it means that the attacker can monitor your activities and freely execute attacks. You will have virtually no visibility or insight into what is happening to your systems and data if unauthorized attachments exist anywhere within your network.

TOPIC 01

HOW TO SUCCEED

01 To be successful, you should **separate your hardware** into three distinct classes. All of your equipment should either be a piece of frontoffice support (such as a laptop or desktop computer), a server used for your infrastructure or WAN equipment such as a router, switch, optical shelf or multiplexer. Each of these three classes of equipment will have varying security requirements and capabilities. Applying a one-size-fits-all approach to secure all of your technology assets is impractical and useless.

02 With your different classes of assets, you should ensure you have **monitoring systems confirming active inventory** for each of those asset types. For example, you can use programs such as Microsoft SCCM or LanSweeper to continually probe your office network and maintain an inventory of the devices it sees. Your WAN equipment should be monitored with a network management solution—whether commercial or home-grown—and use Ping and SNMP probes to continually check the availability of what should be a relatively static environment. In all cases, from a security perspective, you should know right away when any new asset joins your internal IP ranges.

03 Downstream member IPs and assets should be **kept completely separate** from the rest of your organization. You should have a clear understanding of IP addresses and equipment that belongs to—and is under the authority of—the members you serve. Always ensure that subscribers are given dedicated IP space, and never reside on the same subnets as your office computers, server equipment or WAN infrastructure.

Inventory and Control of Software Assets

To create and implement your cybersecurity program, the most critical aspect is to fully understand the hardware assets you are trying to protect. All of our efforts to protect our environment are wasted if there are servers, laptops, switches, routers or wireless access points that we simply do not know about and do not know have to be secured. As a network provider, your hardware is your lifeblood, and ensuring that you have a full and complete understanding of your assets is a requirement for not only providing services to your customers but also making sure you can provide them in a secure and reliable manner.

THREATS

Attackers will try to run or execute malicious software on your devices to gain access to and remove data and information. Such software can additionally operate as malware or ransomware, crippling your organization if it is permitted to exist and execute. Furthermore, attackers may take advantage of vulnerable applications or operating systems to run their malicious actions, bypassing several security protections that would normally find and prevent this dangerous behavior.

TOPIC | 02

HOW TO SUCCEED

01 To begin, you must **keep an accurate software inventory** of each of your three areas of your organization—your desktops, servers and WAN equipment. This software inventory should include both the name of the program and the currently running version and should be directly ascertained by using the hardware inventory in the previous section.

02 To manage your inventory inside your WAN environment, **use scripting or vendor tools** to regularly query your routers, switches and fiberoptic gear. In many cases, this can be performed via SNMP or SSH polling.

03 For your server environment, you should perform daily inventory updates with a management system such as Chef, Puppet, SCCM or other central utility to **keep an accurate listing of both applications and libraries** of active software.

04 When managing your client desktop environment, use PowerShell, SCCM, LanSweeper or other tools to **maintain inventory of all programs installed** on your employee workstations.

05 Once you have an understanding that your authorized software is operating on your authorized hardware, you should **implement software restriction policies or AppLocker** for Windows based computers. Restricting which software can run at the operating-system level can stop malware from executing and spreading, no matter what new vulnerabilities are discovered and exploited.

Continuous Vulnerability Management

All networks and computing environments are dynamic and always evolving. Even with us knowing and enforcing approved software to operate only on approved hardware, attackers will be able to compromise us by taking advantage of vulnerabilities in what we still permit to operate. These vulnerabilities are constantly being discovered and disclosed, and as critical organizations we need to ensure that we are vigilant in not only knowing about these risks as they are discovered but also updating our clients, servers and WAN equipment so they are not compromised.

THREATS

Using vulnerabilities in active software is the primary, traditional way in which attackers are able to hack networks and cause substantial damage and data loss. Several thousand vulnerabilities are disclosed each year, and multiple websites outline not only what these vulnerabilities are and how they work but also give easy-to-use examples and software code on how to turn your well-meaning software into a weapon that is used against you.

TOPIC | 03

HOW TO SUCCEED

01 The first thing your organization must do is **subscribe to vulnerability notifications from your vendors**, including your network equipment providers. These are typically email list subscriptions, and all the members of your IT team or IT security team should have these notices delivered to their mailboxes.

02 As an added layer of protection, your IT and IT security teams should **subscribe to vulnerability notifications from third parties** such as US-CERT, MS-ISAC and CISA. These organizations are constantly keeping track of new vulnerabilities from all technology manufacturers and will help fill in any gaps for any vulnerability notices that your staff may not be directly subscribed to.

03 With this knowledge in place, you need to **run monthly vulnerability scans** on your entire desktop, server and infrastructure IP space. You can use a software product such as Nessus, InsightVM or the Merit CISO Scanner to perform this automated scanning to actively check your systems to see whether they are susceptible to any known vulnerability or misconfiguration. These scans should use a set of administrator credentials so that the scan can obtain the most accurate information.

04 Allow vulnerability data to be **accessible to senior leadership** for decision-making and risk management. Your staff members should be actively reviewing the results of the vulnerability scans and recommending any identified items to senior leadership. This ensures that the responsibility for risk and security is properly communicated and accepted by all appropriate parties and individuals.

Controlled Use of Administrative Privileges

Even if we assume that our computer systems have perfectly managed hardware, software and vulnerabilities, we need to look at the next vector that an attacker would use to harm your organization. The use of privileged or administrator-level accounts is, under normal circumstances, the "normal" way to perform elevated actions in your environment—whether it is granting access to a file share or configuring a router port on your WAN equipment. These special accounts must be strictly regulated and controlled if you wish to retain any secure posture you configure and deploy.

THREATS

Attackers will immediately try to locate and gain administrative-level credentials the moment they have obtained any sort of foothold or access to any of your systems. With administrative credentials, attackers can easily and simply disable or remove any security controls that you have in place, including backups, password security, antivirus and monitoring of any portion of your environment.

TOPIC 04

HOW TO SUCCEED

01 **Remove administrative rights for regular user logins**. This includes local administrative rights for any employees on their local workstations. The vast majority of any performed work will never need this level of access.

02 **Use distinct accounts for administrative work** in your active directory domain. The accounts that can perform elevated actions, such as managing and patching your servers, should be completely different and have different passwords than those set up for daily use. IT support staff should perform their daily work actions using their nonelevated accounts, using "Run As" or "sudo" only as necessary to run the software or functions needed.

03 **Use distinct network accounts for administrative access to WAN equipment**. You should not be using any group or shared logins to manage your network. Each login should be unique so that you can not only manage access as your employees enter and exit your organization but also trace back any work actions to a specific person.

04 **Implement MFA in front of all administrative access**. Whether it will be for workstation, server or WAN equipment management, multifactor authentication should be used to ensure that even if passwords are disclosed, your support staff still retains crucial control of your infrastructure. Using Microsoft MFA, Duo or Yubikeys are all great options to investigate.

05 **Configure an emergency glass-break account** with strong password and alerting. Network outages may result in central authentication not being available. For these scenarios, configure a shared group account on your WAN equipment for emergency access only. The password on this account should be stored in a password management system and be accessible only with proper authorization. To prevent the casual use of this account, configure your monitoring systems to immediately flag, alert and send a group email whenever logins are detected using this shared emergency account.

Secure Configurations

For Hardware and Software for Desktops, Laptops and Servers

With all of the effort placed into creating an accurate and approved software inventory, managing open vulnerabilities on these assets and restricting legitimate elevated access, our next challenge is to ensure that these systems are configured and deployed securely for their first moments on the network. Your computer systems have a myriad of security settings, many of which must be manually enabled and regularly reviewed for effectiveness. Creating and keeping this highly defensible repository of configuration settings will ensure that you do not leave any inadvertent vectors for attackers to compromise you.

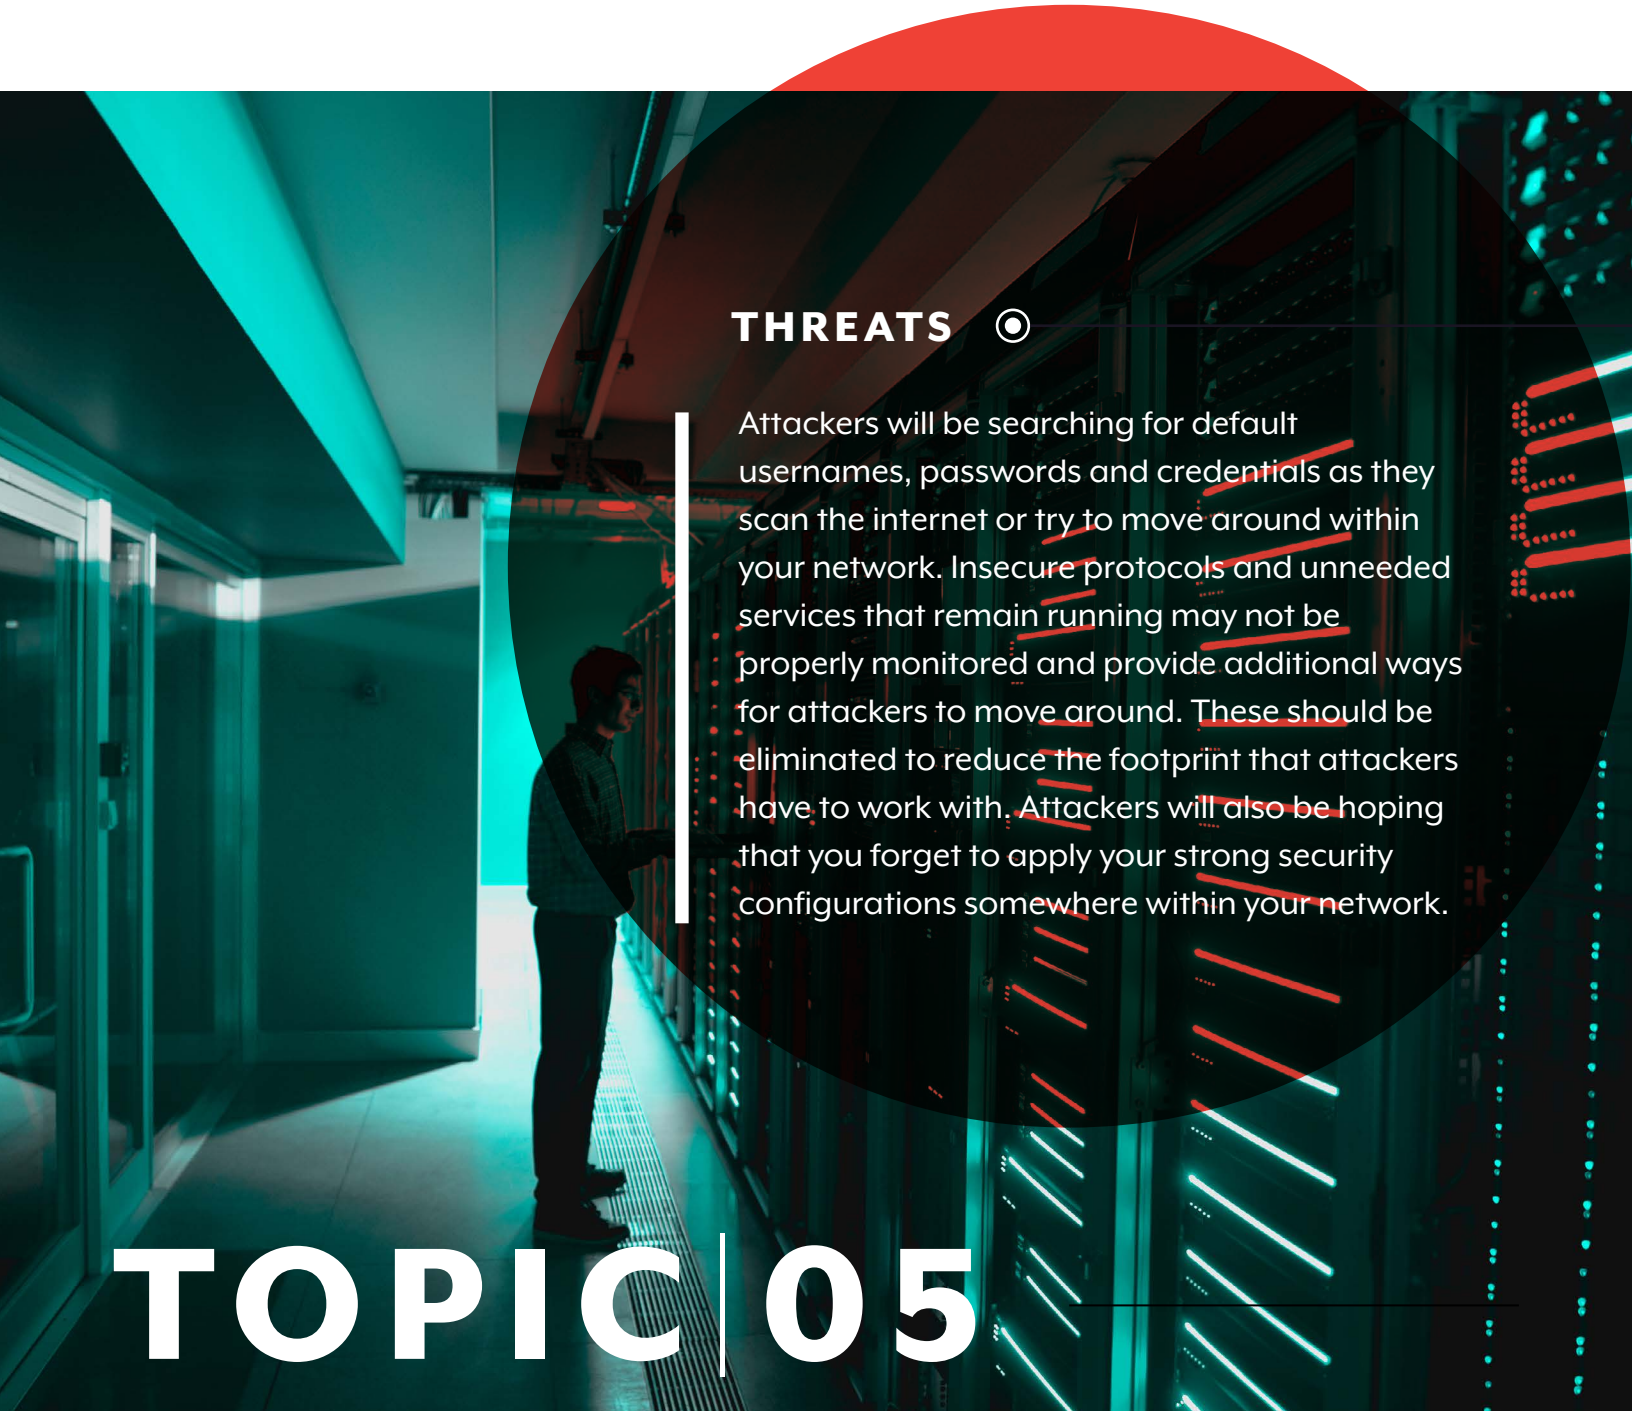

THREATS

Attackers will be searching for default usernames, passwords and credentials as they scan the internet or try to move around within your network. Insecure protocols and unneeded services that remain running may not be properly monitored and provide additional ways for attackers to move around. These should be eliminated to reduce the footprint that attackers have to work with. Attackers will also be hoping that you forget to apply your strong security configurations somewhere within your network.

TOPIC | 05

HOW TO SUCCEED

01 **Have an approved build procedure** for all client and server computing machines. This should include a checklist of required security settings. Recommendations on how to secure these devices are freely available from DISA, the Center for Internet Security and other sources. This process should be reviewed at least twice a year and given to the most junior staff member to test.

02 Software in the trusted build process should be **compared to the active vulnerability lists**, either as part of the vulnerability scanner or vulnerability notices from vendors or other third parties. This should be a point of review both when your team receives vulnerability notices and independently during your regular process for reviewing your build procedure.

03 **Download and store approved copies** of software for your client and server devices in a secure, central location and have their checksums verified and stored in a separate or read-only location. Your build process should only rely on software and configurations that you know and trust.

Maintenance, Monitoring and Analysis of Audit Logs

When it comes to having a great cybersecurity program, it is important to remember that while preventing attacks is great, being able to detect them is a vital prerequisite. Your entire computing environment must have the proper logging and instrumentation to allow your staff to know what actions are happening.

THREATS

For attackers, being seen or having their activities logged is the worst thing that can happen to them. Being able to identify how they are being logged and monitored is typically a high priority, as deleting logs allows them to remain undiscovered for a longer period of time. That allows them to cause substantially more damage and loss to your organization.

TOPIC 06

HOW TO SUCCEED

01 **Log and store all authentication events** within the organization, whether on desktops, servers or WAN equipment. You should always be able to determine when a user account accessed any of your devices and performs any actions that can alter your environment.

02 **Use a central logging host** to store copies of all generated log data. While it is still recommended that devices store copies of their log entries locally on themselves, sending logs over to a separate server ensures that an additional copy exists to work from if attackers are able to compromise any single piece of equipment and delete these local log entries. You can configure several services to accomplish this, such as Syslog-NG or Graylog. This central logging host should have additional layers of security to protect it against any direct attacks.

03 At a minimum, **ensure enough disk space for at least 120 days** worth of log data. Attackers are commonly waiting dormant inside organizations for several months before launching their true attack, and being able to reconstruct events if an attack occurs is crucial to understanding the scope of any breach and what data may have been compromised.

04 **Configure email alerts** for the use of any emergency accounts. All of your IT staff should be immediately alerted if any shared accounts—or other privileged access that cannot be assigned to any one individual—is used. Shared accounts like this are priority targets for attackers, so performing an active announcement whenever these accounts are used is an effective way to not only ensure employees use their own accounts for any access but also alert you if an attacker tries to use them.

Email and Web Browser Protections

Web browsing and email are by far the two most common activities performed by employees on the Internet. As such, these are the primary technical vectors in which attackers and attacks will try to obtain their initial foothold in your environment. For many years, more than 90% of all cyber attacks begin with an employee's interacting with a malicious phishing email. Ensuring that employees are properly protected when using these essential and common services is critical to ensuring that attackers never have even an opportunity to bypass our other security protections and cause us harm.

THREATS

While an organization may think it is safe because it has a firewall that "blocks everything," cyber attackers know effective ways to get their victims—your employees—to come to them. This includes sending emails with malicious URLs or attachments that would have no problem getting through a "block everything" firewall, as the attacker gets the victim to start launching the attack. Web browsers and email clients are not immune to vulnerabilities and are the most common and effective way to bypass other stronger security measures.

TOPIC | 07

HOW TO SUCCEED

01 **Require the use of only one or two web browsers**. These browsers should allow your employees to perform all of their job duties. While Firefox and Chrome are extremely popular, some web applications require either Internet Explorer or Edge. For any supported browsers you select, use your software management system to actively push out software updates for them regularly.

02 **Create a policy surrounding add-ons** or extensions and make sure that you monitor which of these plug-ins are used by your staff members. As part of your software inventory process, you should know which browser add-ons are necessary for work purposes. Restrict the use on any unapproved browser add-ons or extensions, as they are easy ways for attackers to bypass other security measures and run code within this highly used internet application.

03 **Secure your email system** with products to perform both **URL rewriting** and **attachment scanning** and detonation. If you use a popular cloud service for email, such as Office 365 or GMail, you can enable features that closely examine both email attachments and any web links within an email message. For any other systems, a third-party solution such as Proofpoint is a great way to perform these same security checks.

04 **Use a DNS security service** to perform dynamic checks against any malicious URLs whenever an employee communicates to them. This application-agnostic service allows you to prevent communication to known bad or unknown internet domains, no matter which application your employee uses. Cloudflare and Cisco Umbrella provide both free as well as paid options to manage this essential security control.

Malware Defenses

By and large, the goal of attackers is to run unauthorized programs or code on your computer systems so they can extort you, hack you or cause you significant harm or loss. The overarching term for this software is "malware," and it can take many forms. No matter the form, however, the simple fact is that malware somehow needs to execute and run within your computing environment. By taking advantage of several features available within your operating systems, you can ensure that you restrict and prevent the operation of any of these dangerous pieces of software.

THREATS

Attackers have a wide variety of hacking tools to exploit vulnerabilities, download passwords, remove data and any other nefarious actions they care to perform. These malware hacking tools can commonly be repackaged and recompiled to evade common antivirus protections; however, the actions this malware performs cannot be hidden and is commonly what is used to detect and stop execution when it runs.

TOPIC 08

HOW TO SUCCEED

01 **Implement any OS-level protections** against malware execution. Modern operating systems include features that help protect against some of the common vulnerabilities. Enabling these features, such as address space layout randomization (ASLR) and data execution prevention (DEP), typically has zero negative repercussions and can act as an important safeguard if unknown malware is able to run on your computer systems.

02 **Use antivirus or antimalware software** on all desktop systems. At a minimum, enable the built-in Windows Defender module within Microsoft Windows and configure it to perform real-time analysis and periodic file scans. If you have a different product you feel more comfortable with using, enable that instead.

03 **Implement advanced EDR products** on server systems. For your critical server assets, you should enable an endpoint detection and response (EDR) product such as Crowdstrike Falcon or Palo Alto Cortex XDR. These products have advanced features and monitoring, allowing you to easily and quickly track security events on these machines in the event of an attempted hack or breach.

04 No matter which solutions you use for endpoint protection, use their administration consoles and **watch for the detection** of any malicious software within your organization. Many times, attackers will try to run malware, and it is properly blocked from running. By actively watching any antivirus consoles, you can see these attacks take place and respond before the attackers alter their malware so it can bypass your antivirus systems.

Limitation and Control

Of Network Ports, Protocols and Services

One of the ways in which we create a secure and defensible organization is to limit the ways in which an attacker can attach to our systems over the network. Limiting the scope of listening services, especially those services listening by default on our server and WAN equipment, is essential to stopping attackers from launching, pivoting or amplifying their attacks.

THREATS

Attackers will perform reconnaissance either directly or indirectly on the services and applications to both understand how you do business and identify any opportunities to immediately exploit these services for their gain.

TOPIC 09

HOW TO SUCCEED

01 **Use ACLs on WAN equipment** to permit administrative access to originate ONLY from trusted segments of your organization. Inbound SSH or HTTPS connections from outside the IP ranges used by your network engineers should be dropped, and your network should apply BCP38 filtering on all interfaces to eliminate certain risks surrounding IP address spoofing.

02 **Disable BGP on all untrusted interfaces** and require authentication for any BGP or routing protocol peerings. Any routing peerings should be forced to require secure, permissive channels. Where technically feasible, configure both MD5 authentication and TTL Hop Security for BGP peers to continually ensure that routes are being exchanged with trusted partners.

03 Disable, remove or **restrict the use of any insecure protocols**, such as telnet. Any remote administration protocols such as SSH, WMI or Remote Desktop should be limited to originate from your appropriate internal trusted networks only.

04 **Check protocols for old or unsupported versions** of SMB, SSL and TLS—and disable them.

Data Recovery Capabilities

When it comes to the core ways in which we deliver service, nothing could be more critical than ensuring that we have the continual ability to restore service in the event of a significant outage. Being able to respond effectively after a security incident takes place ensures that we can continue operating and servicing our customers.

THREATS

Attackers will try to bring your systems offline, whether intentionally or incidentally during their attacks. For sophisticated attackers, the successful deployment of ransomware will necessitate locating and destroying any critical data backups, as the goal is to get the victim to pay for decrypting their data instead of simply restoring from backup. Even in the case where proper backups are occurring, attackers may try to sabotage the backup infrastructure itself to prevent the successful restoration of your data outside of their control.

TOPIC | 10

HOW TO SUCCEED

01 Keep **daily configuration backups** of all WAN devices within your network. Most of your devices should be scriptable or have the capability to accept remote commands. Use this functionality to keep daily snapshots of the known, good configurations that are trusted to run on your equipment. For certain equipment, you may have the option to have it proactively send a configuration snapshot to a backup server any time you commit a configuration change. In all cases, your configuration backups should go to a minimum of two locations, with one of the locations safely storing data offline in order to be resistant to ransomware attacks.

02 Keep **regular backups for any critical server infrastructure**, including DNS servers, RADIUS/TACACS+, NetFlow or network monitoring systems. Depending on your configuration, you may be able to achieve this either with a system-level backup or just by backing up the configuration and application files for your critical services. To confirm you are covered, test restoration procedures for these services at least annually, and ensure you can restore them within a specified time.

03 Encourage and publish a **single location** for employees to store their data, whether that is a shared network drive or cloud service such as DropBox. Typically, this is where your employees would store contracts, design documents, as-builts, invoices, strategic plans and all manner of business documentation that must remain accessible. Perform an independent backup of this data every day to an offline location to ensure business continuity in the event of a system failure or ransomware attack.

04 Ensure **management accepts the RTO and RPO**, given your backup designs and budget. Your executive team needs to understand, accept and design around two specific timelines governing any information system in your organization: the recovery time objective (RTO) and recovery point objective (RPO). Organizational leadership needs to be clear on the maximum time that a service can be offline; this is the RTO and can be addressed by adding additional layers of redundancy into your equipment, processes and services. Your organization also needs to accept how much data it is willing to lose if its system fails; this is the RPO and is typically addressed by having more frequent and comprehensive backups and ensuring they are stored in diverse locations.

Secure Configuration

For Network Devices, Including Routers, Firewalls and Switches

Your network is your lifeblood and the primary method to deliver service to your customers. Unfortunately, network equipment typically lacks the same high level of security features and protections that managed workstations and servers have. To make matters worse, an attacker having access to your network backbone could permit them unfettered access to the raw traffic of your customers and give them a beachhead into your organization that will be either difficult or impossible to fully remediate in the event of a security breach. Properly controlling access to these crucial network assets is an extremely important endeavor.

THREATS

Attackers see networks as the gateways that allow significant or total access to the rest of their goals. By successfully compromising a switch or router, they can potentially redirect traffic, steal credentials, remove data and position themselves to remain undetected for an extended period of time. Because these components of network gear may not have the same protections against running malware or malicious code that other areas in an organization have, they are also a hotbed for remaining undetected and launching further attacks, either against your organization or a third party.

TOPIC | 11

HOW TO SUCCEED

01 To begin with, **remove all default passwords** from WAN devices and network infrastructure. This includes changing both any privileged or unprivileged user accounts as well as any SNMP communities used for remote monitoring. Any network device that is placed in production or in the field should operate under the assumption that attackers will try to use common passwords against it at any time.

02 **Remove any unneeded services**, daemons and routing protocols from running on your WAN equipment. The running services should be limited to the bare minimum number required for continual operation. This includes protocols or services at any layer of the networking stack. You should be extremely sensitive toward any protocol that broadcasts unnecessary information or capabilities of your network, such as Cisco Discovery Protocol (CDP) or Link Layer Discovery Protocol (LLDP). For any services that would only be occasionally necessary, such as FTP or TFTP for file transfer, keep the service administratively disabled until expressly needed by your operations staff.

03 **Implement a strict ACL** for any traffic destined inward towards any WAN device. Your network equipment should enforce either an access control list (ACL) or other similar technique for any management traffic sent directly to the network device itself. This ACL should be strictly limited to only allow traffic from a segment trusted by your network engineering and operations staff.

04 If possible, **restrict access to an independent VRF network** to increase isolation and prevent unwanted traffic from even attempting to access your network core. This provides an additional layer of protection in separating customer revenue traffic from your internal management traffic and gives you the flexibility of using central advanced firewall protections against traffic destined toward your network infrastructure.

Boundary Defense

When it comes to defending both your organization and those of your members and subscribers, ensuring that proper security boundaries are created and have appropriate security controls is essential to segregating your assets and keeping attackers at bay. Creating effective security boundaries helps identify and isolate assets and focuses security control efforts on these choke points where they may be more cost-effective and useful to your organization.

THREATS

Attacks will constantly attempt to enhance their level of access within your organization. Typically, this happens by pivoting or moving through these common boundary access points. Of specific importance is an organization's primary internet firewall This is typically where deep-packet inspection, antivirus, malware detonation and other technical controls are enforced. If these firewalls are improperly configured or permit traffic in an undesirable manner, attackers can bypass these critical systems and gain a foothold inside the trusted areas of your organization.

TOPIC 12

HOW TO SUCCEED

01 **Implement network-based firewalls** in front of your desktop and server environments while keeping them separate. Your client desktop environment and your server processing environment have different levels of trust, and enforcing those trust boundaries on your network is an important way to stop attacker pivoting and increase the visibility into what devices are communicating with each other inside your organization.

02 Have central and **approved methods for allowing vendors** or support staff to access your infrastructure. Having individuals access your protected assets and systems is a necessary evil. However, standardizing and minimizing this access is key to preventing unauthorized access from accidentally slipping through. Where only periodic external access is essential, require your staff to proactively host a screen-sharing session and monitor the remote access. This way, they have the ability to terminate access and are in full control of the time and method that internal assets are accessed.

03 **Discourage the use of persistent "allow all" firewall rules**. Any firewall rules should be specifically written with source and destination addresses as well as specific ports or protocols. If able, you can also perform time restrictions on firewall rules to better restrict any inbound access through your boundary firewall.

04 Ensure employees **use a VPN with MFA** for any regular or approved inward access into any component of your network. Using multifactor authentication (MFA) is crucial, as it gives you the additional assurance that any access is being performed by your authorized employees and not by attackers using stolen credentials. When your employees do connect through this security device, separate out your network engineering employees and only give their accounts the access needed to monitor and configure your network infrastructure and backbone.

Data Protection

Controlling your institutional and organizational data is an important part of your information security program. Your customers, suppliers and key partners expect you to take the security of their information seriously, imposing necessary precautions to prevent unauthorized access, disclosure or removal of that information. As network providers, we can do several key things to not only proactively ensure that this data does not leave our organization but also engineer several detection mechanisms to hopefully notice and prevent such data leakage in the event that an attacker has some success in breaching your network.

THREATS

One of an attacker's goals is to find and steal information that may prove to be useful or valuable for subsequent action. For example, financial data, contract information, network schematics and strategic plans would all be considered "crown jewels" worthy of identification and duplication. Finding this data and information can be both time-consuming and noisy, so ample opportunity exists for defenders to be able to identify when this activity is occurring.

TOPIC 13

HOW TO SUCCEED

01 **Deploy BitLocker full-disk encryption** to all organization desktop computing devices. Lost or stolen desktops or laptops can be a gigantic liability for your organization, especially if your devices have stored or cached personally identifiable information (PII) stored on them. However, by ensuring that all your organizational devices have their storage fully encrypted when at rest, you can exempt yourself from several consequential aspects of liability, including reporting requirements of the Michigan Identity Theft Protection Act.

02 **Monitor NetFlow** for any large outbound flows of data from your corporate environment. While you typically have no control or concern over data flows to and from your downstream subscribers, you should be alert to any activity coming from your internal servers or desktop environment. Large traffic flows—especially those happening during unexpected times—may indicate either a large removal of data or abnormal reconnaissance scans using your substantial network resources. You can use free software, such as the open-source flow-tools packages, to act as a central repository for NetFlow data, which your routers can then send the information to.

03 **Deploy a honeypot** inside your network to search for problems. To help act as a tripwire for any attackers that may penetrate your defenses, you should set up a network honeypot on one of your trusted internal networks and configure it to alert or email your IT team whenever someone tries to log in or access them. You can download and install a variety of free honeypots, including Cowrie for simulating an SSH server or Dionaea for simulating a variety of other network services.

Controlled Access Based On Need-To-Know

As a network provider, you will have a variety of different employees, contractors and individuals working to help drive your vision and service your community. An important cybersecurity concept is that all of these individuals will have different needs when it comes to viewing and editing different types or classifications of data that you have—and that it is essential to tailor the access given to these people so that they have the minimum amount of access necessary to perform their duties. By properly employing this concept, you can reduce the risk profile at your organization because you are effectively limiting the number of individuals or computers that attackers can target to effectively get what they want.

THREATS

Once an attacker is able to make even small inroads into your organization, some of their next steps will be to identify users or computers that have access to the data or resources the perpetrator wants. For example, the attacker may be interested in the CFO to go after bank account information or the CTO to obtain privileged access into the core network backbone. However, because these individuals may be covered by many layers of security, intelligent attackers may try to compromise other individuals in your organization in the hope that they would inadvertently have access to protected resources.

TOPIC 14

HOW TO SUCCEED

01 Ensure that proper **Active Directory groups** have been created to segregate data and information inside any network drives or shares. Groups should be created for each business function, and central data shares should be provided for each group, with enforced permissions to ensure data remains confidential.

02 **Audit group memberships** every three months to ensure that scope creep has not caused individuals to view or edit data they should not have access to. Many times, employees change roles and we never remove old access from their accounts because "they may need to help out" in the future. This is a dangerous precedent to follow, as many times longtime employees will have access to excessive amounts of files and information that attackers could take advantage of.

03 **Audit all applications**, including cloud or third-party applications, at a minimum of annually to ensure that group access is properly assigned. Any additional computer systems used should have their access groups and access rights checked, including any ERP, GIS, database, network management and monitoring, and telephony systems along with other infrastructure supporting your organization.

04 Ensure that there is a **data owner and responsible party** for every shared folder or application that houses data at your organization. File folders without assigned owners are gigantic risks, as every employee naturally assumes that dealing with any data inside is "someone else's problem." This can cause troves of files and folders that, while containing potentially sensitive or confidential information, are the responsibility of nobody. You can solve this by recording which team or individual is fully responsible for any file folder, share or application that your organization uses. This way, when questions arise as to the necessity of particular files or data records, there is no ambiguity and there is confidence that your data is being fully and properly managed.

Wireless Access Control

Ubiquitous access of Wi-Fi networks is extremely commonplace, with the main use and benefits being the ability of laptops, tablets and personal devices to access internal company resources or offload their data access from cellular networks to a faster alternative. Unfortunately, using organizational Wi-Fi also means that you can have a much more difficult time ensuring that only your own, authorized clients connect to your network. Wireless signals can easily propagate walls and bleed into public spaces, so managing this access properly is critical to maintaining the security of your network.

THREATS

Attackers have long since known that wireless networks can act as an important backdoor into an organization, bypassing many traditional security controls such as firewalls and intrusion detection systems. To this point, numerous high-profile attacks have taken place where attackers have breached networks and stolen critical information while sitting in a car in a nearby parking lot. Network providers have no exemption from this, and unprotected Wi-Fi networks could provide a critical path for attackers to access member data and pivot around inside your entire backbone infrastructure.

TOPIC 15

HOW TO SUCCEED

01 **Use WPA2-Enterprise** for all Wi-Fi access. For maximum security, you should use the strongest protocol set available to your wireless access points. For Wi-Fi networks used by your organizational devices, avoid using either unencrypted "open" networks or older WEP or WPA encryption options. Additionally, avoid WPA2-PSK, as security is only provided by a singular common password that is easily shared and accessible to connected clients.

02 **Use 802.1x** to place wireless clients on appropriate and segregated networks. To achieve security parity with your wired network, you should configure your Wi-Fi systems to perform active segregation for connected clients. Your organizational laptops, tablets and mobile phones should be configured to use a single SSID with your wireless access points using 802.1x responses to place clients on different VLANs or subnets based on their role or purpose.

03 **Use a "guest" Wi-Fi** network for any nonmanaged devices. Guests of your organization expect Wi-Fi access, as do your own employees for their personal devices. A guest SSID can provide this network access, and you can keep your backbone secure by ensuring that clients using your guest network egress on a separate VLAN or on a completely different third-party provider not affected by your backbone routing operations.

04 **Monitor your office regularly** for any rogue or unauthorized access points. Your employees may inadvertently install unauthorized Wi-Fi access points to enhance poor or nonexistent Wi-Fi access for their personal devices. Because these wireless networks may lack proper security protections and may connect directly into your trusted infrastructure, you should perform a walk-through of your office area every few months with a laptop running ViStumbler or InSSIDer or by using a Wi-Fi analyzer to identify any wireless networks available either inside your building or any adjacent public areas.

Account Monitoring and Control

The computer and system accounts for your employees can control the keys to your proverbial kingdom. As such, it is important that we manage the usage and control of these credentials to ensure that they are used only in approved and authorized manners. Of critical importance are groups or accounts that can access the equipment on your network backbone or act as "domain administrator" inside a Windows Active Directory. These accounts, as well as any accounts that have sensitive permissions assigned to them, must be closely monitored and protected to prevent catastrophic damage from affecting your organization.

THREATS

As attackers look to elevate their access and control more areas of your organization and infrastructure, a prime focus will be locating and trying to use usernames and passwords that have significant access to your computers, servers and network backbone. These passwords or related security tokens can, in certain cases, be "scraped" from memory from employee computers if they are frequently used, so minimizing their use drastically reduces the chance of an attack's being successful.

TOPIC 16

HOW TO SUCCEED

01 **Have a standard account provisioning and deprovisioning policy.** As employees transition into, out of and within your organization, you need to have standard procedures that are followed to ensure that inadvertent access is not retained and that you adhere to the concepts of "least privilege." This should be coordinated with your HR and IT groups and be a documented process that can be easily followed. Reports of recent user access changes should be made available to leadership so they can properly manage and accept the risk that arises when assigning and granting privileged access.

02 **Use separate accounts** for any use of administrative rights. While many employees at your organization will have the ability to have elevated or administrator-level access into your network backbone or active directory, those accounts should NOT be used for general, daily computing. Accounts used during web browsing and email checking are highly susceptible to phishing attacks and should have minimal rights to prevent any of these attacks from succeeding. When employees need to perform work on your critical infrastructure, they should use either "sudo" or "Run As" to temporarily elevate their rights for the minimal amount of time necessary to perform their work.

03 **Use email alerts** for use of shared administrative accounts. To enforce that shared accounts and administrative rights should be used as infrequently as possible, configure your logging and audit systems to send email notifications to your IT team or network engineering team whenever such shared accounts are used. These alerts should then be addressed by your staff to confirm that any access was performed by your own employees and not by unauthorized attackers.

Security Awareness and Training

While your backbone, servers and desktop environment can have overlapping technical security controls that are predetermined and highly trusted, they can all be bypassed with a simple bad click by your employees. These attacks come in the form of email, websites and text messages that prey on the emotions and human nature of your staff members. Over 90% of successful hacking attacks—including the large public attacks reported on many media outlets—originate from one of these social engineering attacks. As such, properly training our employees to identify and report these attacks is crucial to having a solid cybersecurity program.

THREATS

By far the most vulnerable aspect of any organization is going to be its employees. Making matters more interesting, electronically attacking your employees may be the only option available if firewalls and intrusion detection systems prevent any sort of direct attack coming outside from the Internet. Attackers can send thousands of convincing yet fraudulent phishing emails per minute, so even having less than 1% of these attacks succeed could prove to be beneficial—as well as dangerous for the victim organization.

TOPIC | 17

HOW TO SUCCEED

01 **Perform regular training to your employees** on proper computing practices. Your employees should be exposed to cybersecurity topics at least on a monthly basis, with an emphasis on how to identify and report potential cybersecurity attacks or incidents. This can be through a selection of video-based training through a service such as KnowBe4 or SANS, free periodic newsletters, regular team meetings or email updates sent by your IT and cybersecurity leaders on the latest topics surrounding staying safe online and cybersecurity.

02 **Perform simulated email phishing** campaigns to help target your training. The most effective way to make sure your employees can spot social engineering attacks is to regularly expose them to realistic attack examples. Not only will this continually reinforce the lessons taught in your security training, but it will also allow you to identify particular kinds of attacks that could be more successful. Free tools such as GoPhish or other commercial options can get you operational in just a few minutes. Whatever you choose, keeping an assessment program like this positive and upbeat is crucial for employee morale and engagement.

03 **Report on education success** to both employees and management. Cybersecurity is everyone's responsibility, and part of that is to ensure everyone has feedback to know how effective efforts are toward your goals. Senior leadership should be made aware of both the successes and failures of training and assessment, with staff members receiving positive reinforcement for when they identify potential attacks or help mentor other employees in safe computing practices. Maintaining long-term trends of success will also keep your employees and leadership engaged and committed to keeping your overall network secure.

Application Software Security

As network providers, we will be developing our own software, scripts and programs to help us run our organization and deliver service to our customers. Whether it is gathering bandwidth or flow data for billing or interface and device monitoring for our NOC, we will be using custom code tailored specifically for our business needs. As we do this, we need to ensure that the software we develop does not include vulnerabilities or backdoors that could cause us significant inadvertent harm.

THREATS

Your applications expose both your data and your infrastructure, and attackers can take advantage of that if inexperienced development teams introduce vulnerabilities that can be exploited. Data disclosure on error pages, directory traversal, file inclusion, lack of input sanitization and SQL injection are all common aspects that attackers can use to access unauthorized subscriber data, list out your data structure or potentially even run code of their own choosing on your systems.

TOPIC 18

HOW TO SUCCEED

01 **Use proper coding standards** using the OWASP Top 10 as a model for proper web development security. While developers can make many common mistakes when creating software code, in most cases this simply stems from not fully knowing how software or web applications can be misused or hijacked. Imbue your development team with a secure mindset by providing training and resources on how to code securely. In addition, have any developed code be reviewed by another developer to ensure that your coding security guidelines are being followed properly.

02 **Use a password manager** for storing credentials that are used by your applications. Specifically, using a privileged access management solution such as Thycotic can give your code the flexibility to "check out" credentials, tokens or passwords through an interactive API. By doing this, you can cycle your service account and token passwords regularly without worrying about breaking your key applications.

03 **Never check passwords or credentials into a code repository**. Especially if your team is using a solution such as GitHub or Bitbucket to perform code development, ensure that static passwords are never included in checked-in files and are only made available in local files. Any checked-in files in your development system will be stored indefinitely, meaning that any credentials accidentally submitted will never be able to be fully removed or erased.

04 **Run a vulnerability scan against developed applications** before moving them into production. As part of your quality assurance and review process, you should use your vulnerability scanner to perform a final validation that no significant errors exist for an attacker to discover and exploit. Besides your regular vulnerability scanner, you can use either a commercial option or a free and open-source solution such as Nikto.

Incident Response and Management

While we place great efforts into ensuring that our networks remain secure, it is only a matter of time before a perceived hack or breach is identified and needs to be addressed. Besides ensuring that you have effective mechanisms for reporting suspected security incidents, it is important to have structure and policy in place so that such reports can be quickly and effectively triaged. With attackers typically being able to hide inside organizations for around 180 days before being detected, having a strong detection and response plan will ensure that you can minimize any risk and loss as a result of an attack.

THREATS

Attackers know that no matter how successful their initial attacks may be, they will be discovered at some point in time. That being said, if an organization lacks an effective incident response capability, attackers can remain in a network longer, giving them additional crucial time to find your data, compromise additional systems and ultimately achieve their goals.

TOPIC 19

HOW TO SUCCEED

01 **Have an incident response plan with appropriate roles** and assigned individuals. An incident response plan does not have to be complicated or technical but simply outline the broad steps used when responding to a cybersecurity incident, whether big or small. Your plan should designate staff to receive and review security incidents as well as define different classification levels so that minor security incidents (such as a phishing message) have priorities and resources that are different from those involved in a large-scale breach (such as losing your entire customer database).

02 **Ensure your plan is accessible** to any employee at your organization who would be playing a role in managing an incident or disaster. Your plan should be in both electronic and hard-copy formats and generally available to any primary or secondary support staff listed inside the incident response plan. Because of the unpredictable nature of security incidents or disasters, having this plan be available via multiple modalities—including via cloud-based storage—is a great way to ensure that you can quickly respond, no matter the nature of the security incident.

03 **Update your plan annually** or during significant infrastructure changes. Your technical and leadership teams should set time aside each year to ensure that any technology, staffing or policy changes are properly reflected in your ability to respond to a security incident. An outdated plan that references obsolete technology or nonexistent staff will hinder any effective response to a cybersecurity incident, causing additional potential loss.

04 **Perform a tabletop exercise annually** to test your plan against realistic scenarios. Once your incident response plan is updated with the most appropriate procedures and information, your leadership team should simulate some potential attacks and ensure that they are properly addressed by the contents of your plan. Examples could include a ransomware attack involving finance systems or a water main break that destroys portions of your network's infrastructure backend.

Penetration Tests and Red Team Exercises

With all of the above security controls and precautions, the final piece to ensuring that your network remains secure is to perform periodic testing from the perspective and mindset of an actual attacker. Performing periodic assessments in this manner validates all of the technical, administrative and staffing resources that you have invested in securing your network and can help prioritize additional security projects and initiatives to continually reduce your risk and create the most secure network that your communities can trust.

THREATS

Attackers can be creative in identifying vulnerabilities and finding unique chains of actions that can result in a successful hack or compromise. Experience has shown that attackers will be tenacious and use complex combinations of attack vectors to discover your network, exploit vulnerabilities and embed themselves inside your infrastructure.

TOPIC | 20

HOW TO SUCCEED

01 Have your IT or security staff **perform dedicated threat hunting** at least once per quarter. This consists of setting aside time for your technical teams to search through your client, server and WAN equipment to confirm that systems are configured correctly and that no suspicious files, processes, network connections or user accounts are either in active use or lying dormant. Threat-hunting teams can use an endpoint detection response product such as OpenEDR to quickly scan computer systems for potential indicators of compromise (IoCs).

02 Have an outside firm **perform a penetration test** at least once every other year. Third-party firms are a great way to get an unbiased and informed view on how all of the pieces of your cybersecurity program operate and integrate. Additionally, they can identify any weak spots or misconfigurations that may not be readily apparent to your own staff. External penetration testing reports can be used to prioritize and fund large-scale security improvements which are necessary to remain competitive in delivering service.

03 **Reward your staff** for identifying security problems and working in a secure mindset. Cybersecurity is everyone's responsibility, and your staff members will always have the best perspective on identifying any suspicious activity in your systems. They have the daily exposure to what the response times and expectations are for working with your backbone and information systems, and encouraging them to not hesitate to report potential issues is a key to success. Always congratulate your staff when they report incidents, pointing out and confirming any "red flags" that caused them to be suspicious. Using token incentives to reward reporting phishing messages or other attacks is an example of giving more ownership of security problems to your employees—the ultimate goal in fostering a security mindset throughout your organization.

About Merit's CISO Kevin Hayes

Kevin Hayes is the Chief Information Security Officer at Merit Network, Inc. In this role, Kevin is responsible for the management of IT security controls and products, responding to information security incidents big and small, and providing security policy and strategy guidance to Merit members, system administrators, and management. No stranger to the nonprofit world, Kevin has previously worked at Wayne State University, creating and directing the team of cyber professionals dedicated to keeping the organization secure. Kevin holds both CISSP and CISM certifications, is a member of the Governor's Cyber Civilian Corps, and has been heard talking about security issues from time to time on the Detroit television stations WDIV-TV and WXYZ-TV, as well as the Detroit public radio station WDET-FM.

CONNECT WITH KEVIN:

LinkedIn | Twitter