Toward a Zero Trust Architecture Implementation in a University Environment

Erik Dean
Shane Fonyi
Lieutenant Colonel Christopher Morrell, Ph.D.
Dr. Michael Lanham
Colonel Edward Teague, Ph.D.

ABSTRACT

The core concepts of Zero Trust Architecture have existed since the Jericho Forum in 1994 and have served as the goal of cyber security specialists for many years. Zero Trust Networks and Architectures are extremely appealing to institutions of higher learning because they offer the flexibility to support research and learning while protecting resources with different protection levels, depending on the sensitivity of the resource. This paper investigates how other universities can employ the Zero Trust Architectures using the West Point model.

INTRODUCTION

Traditional network architectures focus on a static defensive perimeter augmented by multiple static layers of additional security which are more than sufficient when resources within the perimeter remain in fixed locations with a user population located within the same perimeter. With more mobile users it does not work, especially as cloud computing becomes more prevalent. These new circumstances require

Erik Dean is the Chief of Information Technology and Security Operations at the Army Cyber Institute. He conducts research in emerging technology implementation and integration. He has previously worked for the Office of Economic and Manpower Analysis and the Information Technology Operations Center. Mr. Dean has a background in Computer Science, Criminal Justice, and Network Technologies.

organizations to adapt policies and procedures to fit them. Organizational control and security over data stored in the cloud and accessible from the internet is generally more difficult on networks with only on-premises controls. Once data leave a physical location, networks with traditional policies and procedures are often unable to secure that data from unauthorized party access. The same goes for devices that connect back to the network using legitimate credentials without sufficient scrutiny of that connecting device. Exclusive focus on the user and security of the credentials exposes the device to compromise even with necessary protections in place, such as multi-factor authentication. Zero Trust Architecture (ZTA) is not only about technical controls to prevent unauthorized access, but also about policies that promote a more secure and mobile workforce. The concept of defense in depth is the main principle in focus for this architecture type, but now with a greater focus on endpoints outside the network perimeter.

With ZTA, there is an assumption that there is no inherent trust between two assets. All connections are scrutinized as if they were previously unknown. Authentication and authorization are separate functions that must occur before a session can be established with an enterprise resource.[1] When a user attempts to connect to a resource from any device or network, the user must be authenticated, the device must be trusted, the resource must be verified, and finally the authorization for access to the resource must be validated. Only after the Zero Trust workflow is completed, can a session be allowed, and the user given access to the data. The concept differs from traditional networks that automatically trust all connections within the internal network enclave without scrutinizing the endpoints making the connections. If the network traffic is allowed, then the session will be established in most instances. This will require organizations to confirm that their controls and policies currently address these topics and can adapt to the changing environments.

Shane Fonyi is a Cybersecurity Analyst for The United States Military Academy at West Point. He is responsible for responding to and preventing cyber incidents at the Academy as well as the engineering and maintenance of the IT security systems currently in place. Prior to that, he was a Cyber Research Integrator for the Army Cyber Institute at West Point and a Security Engineer at The University of Kansas (KU). He has been a speaker at several IT and information security conferences including the International Conference on Cyber Conflict, BSides KC, and the Conference on Higher Education in Kansas for work involved in 5G security, IoT research, and security awareness. He currently holds a Bachelor of Science in Electrical Engineering from KU as well as several industry certifications including the CISSP, CySA+, GMON, GCFA, and SSCP. He is also an NYU Cyber Fellow at the NYU Tandon School of Engineering.

Some organizations have implemented bring-your-own-device (BYOD) programs, but many of those still have major organizational security concerns, and few as yet have solutions.[1] This paper addresses the National Institute of Standards and Technology (NIST) recommendations for implementing Zero Trust Architectures,[2] from both a policy and a technical perspective, and how the NIST recommendations might apply to University networks that track the West Point network as an example.

BACKGROUND

In early 2018, DoD and US Military Academy (USMA) leadership determined that the network security paradigm applied to traditional DoD networks was insufficient to allow USMA cadets and faculty to foster the kind of academic rigor required of one of the nation's top educational institutions. The decision was made to transition the USMA network and architecture to a design more closely aligned with those found at other academic institutions, to include a Zero Trust Architecture that provides an equivalent level of security mandated by DoD while ensuring the flexibility demanded by academic research and education.

DATA AND COMPUTE AS RESOURCES

As an institution of higher learning, the U.S. Military Academy has a broad range of technological and data resources that were considered for inclusion into the Zero Trust Architecture. This breadth of data and resources is compounded by the fact that USMA is also a DoD asset and has other resources not common to other universities. These resources were considered for inclusion based on their access to the West Point Research and Education Network (WREN). In this case, the selection criteria were simple in that the resources were included in the assessment only if the resource in question can be accessed by WREN users or utilize the WREN for network transport.

Lieutenant Colonel Christopher Morrell is an Associate Professor and the Director of the Cyber Science Program within the Electrical Engineering and Computer Science Department, U.S. Military Academy, West Point, NY. He holds a Ph.D. from Virginia Tech where he studied Moving Target Defense using IPv6. His research interests include network security and optimization, network management, and mobile device software development. Email: christopher.morrell@westpoint.edu.

After considering all devices and data sources that use WREN to transport or process, the following categories were used as resource groups to determine access levels during the dynamic access and authorization steps discussed below, as follows:

1) Personally-owned devices with no health, security configuration, or compliance checking.

2) Enterprise-owned devices and systems that do not support network-based authentication. These devices require network transport and some level of WREN resource connectivity, but configuration and compliance cannot be checked automatically. This category of resources and devices may present a higher risk to other WREN resources.

3) Devices able to (a) perform automated health and security policy compliance checks, (b) be integrated with the device management solution chosen, and (c) perform challenge/response authentication at the network level.

4) Systems, devices, or applications that contain or process Personal Health Information (PHI) or medium-to-large volumes of Personally Identifiable Information (PII), whether or not able to perform challenge/response authentication or report device health/configuration information. These require the highest level of protection.

These categories apply to data contained both within Information Systems (ISs) and the devices.

COMMUNICATIONS SECURED BY NETWORK LOCATION

By design, WREN resources generally must be accessible from anywhere on the planet. As with most contemporary top-tier universities, several West Point cadets participate in immersive study programs, and travel abroad, as is true with West Point staff and faculty, and all require continuous access to WREN resources. Frequent travel is common for most universities,

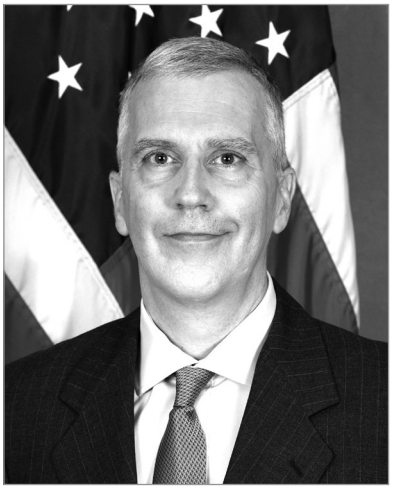

Dr. Michael Lanham is an Associate Professor with the Department of Electrical Engineering and Computer Science (EECS) and Chief Information Security Officer (CISO) of the United States Military Academy (USMA). Michael served 27 years in the Army, culminating as an Academy Professor at USMA with EECS. Prior to that he was a FA53 - Information Systems Management officer since 2003. LTC Lanham was commissioned into the Infantry in 1992 from North Carolina State University. He has served in numerous deployments to Macedonia, Bosnia-Herzegovina, Sierra Leone, Liberia, and Kuwait. His military assignments included duty with 2-15 Infantry Regiment, 3rd Infantry Division (Mechanized) (Schweinfurt, Germany) and Special Operations Command Europe (SOCEUR) (Stuttgart, Germany) as well as with the 1st Brigade and D Co/1-327 Infantry Regiment, 101st Airborne Division (Air Assault) (Fort Campbell, Kentucky). He has also served as faculty at USMA, in various staff positions with USSTRATCOM, Joint Functional Component Command (JFCC)-Integrated Missile Defense (IMD), JFCC-Network Warfare (JFCC-NW), USARCENT, USASMDC/ARSTRAT/ARFORCYBER, and ARCYBER.

which creates challenges for restricting communications based on location and can significantly degrade end-user services. The NIST recommendation therefore would seriously impact user's ability to do their jobs.

Due to these factors, most WREN resources are not restricted by location. The security tools embedded in cloud computing platforms like Google Workspace and Microsoft Office 365 enable this. WREN's cyber security staff leverage the advanced threat identification and mitigation tools these platforms have, in order to compensate for the inability to restrict by geographic location or network location.

While generally unrestricted, there are network controls and restricted access within the fourth category of resources above, which apply only to the local West Point network enclave and do not need external location access. These resources are restricted to on-premises users whose devices meet all requirements for authentication and device health attestation and validation, facilitated by multiple mechanisms such as geographic location via the Company Portal device management platform, client-provided network address (also identified through the Company Portal software), and, finally, the network group to which the device and user have been assigned. If the user and device are trusted, meet all compliance criteria, and are either geographically or logically, though some remote access mechanism such as virtual private networking (VPN), located at West Point, they can access those restricted resources.

WREN takes this requirement further than traditional networks with Software Defined Network (SDN) and the capabilities it provides.

ACCESS GRANTED ON PER-SESSION BASIS

Authorized users with personal and enterprise-owned devices gain access to WREN resources primarily through web interfaces or web-based portals.

Colonel Ed Teague is the Chief Information Officer (CIO) and G6 at the United States Military Academy (USMA) at West Point, NY. Ed, a 1995 USMA graduate, commissioned in the U.S. Army Aviation Branch and flew the OH-58 A/C, OH-58D, and AH-1 serving in Schofield Barracks HI, Fort Drum, NY, the Republic of Korea, Arlington, VA, and Afghanistan. Ed also served in the Operation Research/Systems Analysis Branch, as an assistant professor, and as a program director in the USMA Department of Systems Engineering. Ed is currently an Academy Professor at USMA with a BS in Mechanical Engineering, MS in Operations Research from the University of Texas at Austin, and a Ph.D. in Systems Engineering from the University of Virginia.

This service access implementation paradigm allows for standardized service implementation and access for all clients, reduced developer workload, more focus on specific service entry points, and fewer service entry points that need to be monitored. Using web interfaces as a standardized access method allows standards and compliant session handling to be offloaded onto applications that implement the HTTP and HTTPS web-based protocols, such as OneDrive, SharePoint, and Office.com.

Standardizing access protocols also ensures that authentication and authorization occur through each user/service interaction and are implemented and enforced through well-defined protocol handlers. For the WREN, this has been implemented by centralizing resources access through the Microsoft Office 365 cloud-based platform. By leveraging Microsoft's authentication controls and device configuration management tools, WREN enforces correct authentication and authorization and can enforce device health controls for required resources at a per-session level. Per-session authentication and authorization (A&A) is automatically provided to all enterprise service and network architecture. Any services below the enterprise level (e.g., Academic Departments, Research teams, etc.), are not guaranteed session authentication and authorization as they exceed what the enterprise services provide.

While Office 365 provides robust session handling capabilities, the zero-trust architecture extends centralized authentication and authorization capability solely to services that understand Security Assertion Markup Language (SAML). Any services that do not support SAML must implement this level of authentication and authorization in other ways which, sometimes, do not exist for smaller or legacy applications and software packages.

ACCESS GRANTED BY DYNAMIC POLICIES

As discussed in the next section, by leveraging standards-based protocols, WREN heavily relies on HTTP and HTTPS for session management and handling. Authentication and authorization are handled through these protocols for most applications, and other checks are in place that are enforced depending on the resource being accessed. These policy checks occur through a variety of factors used to determine authorization to access a specific resource.

The first decision criterion used to access the network itself is a comply-to-connect mechanism. Devices must be known to the enterprise architecture through the mobile device management software or through the Microsoft Azure Active Directory domain and must support IEEE 802.1x network-based authentication. For devices unable to support this requirement, other options for device registration and accounting can facilitate the decision-making criteria as to whether a device can access another WREN resource. This check requires the device to be locally resident to a WREN network enclave. Implementation of a purely software defined network (SDN), as discussed in the next section, is governed by technical measures for these local network connections.

The second decision point is device health attestation. Devices must comply with several device health requirements in order to be considered healthy enough to access WREN resources. This compliance is managed through multiple means, including the mobile device management components of Office 365 and Azure Active Directory. Multiple factors are used to generate a health score. Each resource category listed above requires a minimum score. Device health is routinely monitored and devices that fall out of compliance are automatically disabled until corrected.

The third rule, which is related to the first, is geographic location. Depending on the resource type requested, geographic location could be a factor. Some resources are configured with access restricted to a certain geographic area and are hence unavailable to users outside those areas.

ALL DEVICES ARE IN THE MOST SECURE STATE PRACTICABLE

All network engineers aspire to have all connected devices in the most secure state possible, but this is an extremely limiting goal in a research or academic setting. WREN's design into multiple security types has resulted in numerous implementations to meet this requirement. The first category of resources, personally owned and completely unmanaged devices, do not require any specific security settings, because they access only public resources and lack any ability to interact with other WREN resources.

The second resource category is enterprise-owned devices which support teaching or research for a limited time but do not support network-based authentication and authorization or device health attestation. WREN's infrastructure requires no compliance model for these

devices beyond industry best practice, NIST controls, and DoD policies where applicable. While this may seem counterproductive under this model, the devices are required for a limited period and then removed from WREN.

The third resource category includes devices that support device health attestation, network compliance checks, and are integrated with the enterprise device management solution. These devices are typically mobile computing platforms (tablets and laptops) or desktops and use a common security baseline that applies a minimal number of enterprise-level controls. These devices support teaching and research and, when stringent security policies are applied, have significantly degraded performance in routine computer use (e.g., emailing), but also as a research or education computer. Functions such as code compilation, tools that are graphic-processor intensive, or need for precision response times are greatly impacted by most security policies. The security policy load is reduced on these devices, yet many enterprise policies are enforced and monitored that ensure core device health (updates, current antivirus and anti-malware, device behavior, etc.). This category contains both West Point-furnished devices and personally owned devices that require access to WREN resources. Users agree to an acceptable use and management policy agreement and allow application of WREN security policies to these devices. The ZTA implementation technology allows collection of threat indicators from both types of devices to ensure WREN data security, and to provide threat metrics and indicators to the WREN Cybersecurity enterprise.

The fourth, high-risk resource category refers to data protected by multiple regulations which, if compromised, would seriously and adversely impact the user population. Because these computing resources do not directly map to the teaching or research mission and are integral to USMA's core business, they are secured using the most stringent set of security controls.

AUTHENTICATION & AUTHORIZATION ARE DYNAMIC

One of the most difficult tasks for implementing any enterprise service delivery is providing real-time evaluation of user access to a resource. What if, post-authentication, a user's risk posture is reduced? In traditional networks, this time window provides a vector for malware or insider threat actors to access resources they may otherwise no longer be allowed to access.

WREN solves this problem through Microsoft's implementation of the Continuous Access Evaluation Protocol (CAEP),[3] which features a re-evaluation mechanism for each resource request, thus allowing resource administrator control of access to each resource on a per-request basis. This ensures that once a user's access is terminated, the time lapse between the user access revocation and access denial is limited to the time it takes to communicate between the centralized user access control and the resource provider service.

Resource access is governed by WREN's Comply to Connect (C2C) policies, and basic user role-based access (RBAC) controls. Higher-sensitivity resources, such as Privacy Act or educational record data, require the user to meet more stringent configuration polices such as coming from a West Point-issued device, within the physical network enclave of West Point, and having a user account with a low risk rating. Resources with a lower sensitivity level are accessible by users with a wider range of devices that include personally owned but Azure Active Directory-registered devices, a smaller set of security requirements, are geographically distributed, and have slightly higher risk profiles. High risk profile users have access to the smallest number of resources through the fewest number of devices. As a user's risk status rises, their ability to access resources are commensurately reduced.

ENTERPRISE COLLECTS AS MUCH NETWORK AND COMMUNICATIONS INFORMATION AS PRACTICABLE, INCORPORATING CONTINUOUS IMPROVEMENTS

WREN is designed to capture all forms of data, not only for network security but also for network optimization and performance tuning. The data capture draws from myriad sources and can be expanded to ingest nearly any type of data. WREN utilizes the capabilities of Microsoft Sentinel[4] to provide event management and automated response. Originally designed as a Security Information Event Management (SIEM)/Security Orchestration Automated Response (SOAR) platform, WREN leverages Sentinel's robust scripting capabilities as well as native integration with the Microsoft PowerBI platform to yield performance metrics both for the Microsoft Office 365 platform and for local enclave performance. This monitoring occurs through the native logging capability built into the enterprise network devices, connectors to the Sentinel platform, and automated analysis capabilities available once data are stored. Sentinel's integration capability also allows for the ingestion of external security and performance data through protocols such as TAXII and Microsoft, and other third-party threat data. This integration of external threat data along with the powerful scripting language supported within Sentinel also allows the platform to automate many response actions to event correlations which may or may not be an active threat in the network. This also allows Cyber Defenders on WREN to implement threat identification and mitigation capabilities more advanced than those existing on traditional networks.

Using data analysis, WREN's planning team can identify additional capabilities needed to expand existing and future projected capability. By monitoring network performance through Sentinel logging, Cisco DNA, and SolarWinds, network and security staff can identify service disruption due to misconfigurations, infrastructure failure, or unexpected load on key devices. Once flagged, the WREN Network Operations team corrects these service interruptions. Data trend analysis forecasts future bottlenecks or infrastructure challenges that may otherwise be unobservable. This trend analysis is critical in performance prediction and helps identify infrastructure changes, additions, or reconfigurations that can be planned as part of a long-term strategic lifecycle plan.

Collected data, while extensive, is only used to improve WREN's connectivity, throughput, service delivery, and the network's security posture. The collected data are primarily instrumentation data from network devices, performance metrics from cloud-based virtual machines, and Microsoft Office 365 performance metrics. WREN captures some user data, but it does not collect or analyze user-level data by design.

CONCLUSION

The current WREN network implementation is an imperfect model of the Zero Trust Architecture, but it can serve as a road map for higher education institutions that are designing or modifying their networks. West Point will continue to pursue a true Zero Trust Architecture for the WREN and continue to implement technologies that provide a rapid fielding capability for innovative ideas in the educational space and provide a safe, secure, and stable computing environment that leverages both security and optimization at every level found in the ZTA concept.

ACKNOWLEDGEMENTS

This research was enabled by the Department of Defense, the Department of the Army, and West Point senior leadership.

NOTES

1. M. Astani, K. Ready, and M. Tessema, "BYOD issues and strategies in organizations," *Issues In Information Systems*, 2013.
2. S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero trust architecture," *NIST Special Publication 800-207*, 2020.
3. S. Cuff, "Azure ad security enforcement with continuous access evaluation," *URL https://bit.ly/2PtKwrB, 2020*.
4. M. Corp, "What is azure sentinel?" *URL https://docs.microsoft.com/en-us/azure/sentinel/overview,* 2019.