

The Road to Zero Trust (Security)

Co-authors: Kurt DelBene, Milo Medin, Richard Murray

Staff Lead: Schuyler Moore

9 July 2019

BLUF: Zero Trust Architecture (ZTA) has the ability to fundamentally change the effectiveness of security and data sharing across DoD networks. From a security perspective, ZTA can better track and block external attackers, while limiting security breaches resulting from internal human error. From a data sharing perspective, ZTA can better manage rules of access for users and devices across DoD to facilitate secure sharing, from the enterprise center to the tactical edge. Furthermore, the network design and flexibility of ZTA will help DoD more rapidly adopt and implement critical network technologies and enablers, ranging from cloud computing to artificial intelligence and machine learning.

DoD cybersecurity is at a critical juncture. Its networks are growing in size and complexity, requiring massive amounts of rapid data transfer to maintain situational awareness on the digital and physical battlefield. This expansion is stretching existing cybersecurity apparatuses to their breaking point, as an ever-growing number of users and endpoints increases the attack surface of the network. This challenge is not unique to DoD - the commercial sector is facing the same challenges, and their networks are constantly building new connections to a broad range of other networks that drive new vulnerabilities. As a result, public and private sector are re-assessing the current method of “perimeter” security and are considering new methods. One such method is “zero trust,” which could drive a step-change in security improvement across commercial and DoD networks.

Zero Trust Architecture (ZTA) can significantly offset vulnerabilities and threats across DoD networks by creating discrete, granular access rules for specific applications and services within a network. Some of the most severe cases of network breaches could have been prevented using basic zero trust principles - for example, had ZTA access rules been applied to Edward Snowden, he would have been unable to obtain the broad range of documents that he released to the public. Instead, he was given “system administrator” privileges within the NSA network, which provided him blanket access to resources and files. This method of blind trust in users and devices inside the perimeter the network is not sustainable, and will continue to put national security information and operations at risk until it is resolved.

Cybersecurity Current State: Perimeter Security

The rapid evolution and expansion of the digital world has come with a growing number of increasingly sophisticated cyber threats. But despite these developments, cybersecurity practices have only made modest improvements over the past few decades. Early users relied on anti-virus scanning to detect and remove viruses from individual devices, and these practices ultimately evolved to include endpoint protection, threat detection, and response to defend the

broader network. These early networks had a limited number of end points and users and subsequently could rely on “perimeter” security, which emphasizes guarding the entry/exit points to the network by checking the user’s identity and the data packets that come in and out.

However, this approach rapidly becomes strained as networks expand to include a vast number of endpoints, with users requiring access from various locations. Additionally, adversaries continue to find creative methods of getting around perimeter security, such as social engineering attacks that manipulate users into giving away their credentials. This network expansion and adversary creativity requires increasing numbers of firewalls with complex inspection specifications, which is cost-intensive with diminishing returns. Perimeter security rapidly devolves into a game of “whack-a-mole,” where firewalls must constantly be adjusted to account for an expanding set of authorized entrants into the network and acceptable traffic in and out of the network. These firewalls will frequently duplicate efforts, creating burdensome and costly infrastructure that is largely ineffective.

DoD networks face this strain, but on a larger scale and with higher stakes than private sector. As the battlefield continues to rapidly evolve and adversaries close the capability gap with DoD, DoD will require elevated situational awareness to enhance decision-making, and this will require massive data sharing in real time across a wide, diverse set of systems and platforms. Various technology advancements will increase DoD capabilities in this direction: for example, the shift from 4G to 5G will drastically impact the future of global communication networks and fundamentally change the environment in which DoD operates. 5G will enable a higher volume of data shared between more systems and platforms at a faster rate, and has the ability to enhance DoD decision-making and strategic capabilities, from the enterprise network to the tactical edge of the battlefield. 5G will increase DoD’s ability to link multiple systems into a broader network while sharing information in real time, improving communication across Services, geographies, and domains while developing a common picture of the battlefield to improve situational awareness.

This rapid data transfer is critical to the development of future DoD capabilities and enhanced situational awareness. Artificial intelligence and machine learning applied to DoD systems will depend on that flow of data, which will in turn enable a host of new technologies and missions, from resilient satellite constellations to mesh networks. However, each of the added endpoints and data transfers within this expanded network creates a new opportunity for adversaries to target data and operational capability on DoD networks. As the attack surface expands, perimeter security will increasingly become overwhelmed, allowing more unauthorized users to slip into the network. Just as these technologies generate more data traffic, they will also enable the volume and speed of threats to increase, making it difficult for any perimeter security system to monitor and manage those threats. Without a change in cybersecurity strategy, DoD runs the risk of compromising its data, networks and operations.

The Next Generation of Cybersecurity: Zero Trust

Shift from Perimeter to Zero Trust

Compare the experience of living in a house versus living in an apartment building. In your house, there are only a handful of entrances and a handful of familiar people with keys to those entrances. For these reasons, you probably don't lock all the doors inside your house because you have faith in the "perimeter security" (in this case, locked doors leading to the outside with a select list of people with access). In an apartment building, there are many more points of entry and a longer list of people with access, which decreases your familiarity with other access holders and increases the risk of unauthorized access. For this reason, you likely lock the door to your apartment instead of just relying on the perimeter security of the apartment building, because you have less certainty that every person in the building has authorization to be there.

Early networks may have looked like houses, but they are increasingly moving in the direction of apartment buildings, with constant internal traffic offering potential access to each unit. Networks have become widely dispersed across a complex web of connections to outside servers and other networks, with larger numbers of "tenants" and a growing number of entry/exit points. As noted earlier, perimeter security has become increasingly expensive as it requires more firewalls with complex filtering capabilities to cover the expanding attack surface of a network, all while the security of the network continues to decrease as regular usage creates new vulnerabilities. For this reason, it is useful to assume that the network is compromised and take a more targeted approach to security.

What is Zero Trust?

ZTA moves away from the perimeter approach and instead assumes, as the name would suggest, zero trust in the network itself. By assuming that the network is compromised, security can take a more nuanced approach by guarding access to the resources within the network and building strong authentication and authorization standards to allow specific access based on user- and device-specific attributes. ZTA operates on a "least-privilege access" model by only granting users and devices access to the applications, services and data that are absolutely necessary for their role within an organization. By using "role" as a centerpiece for determining access, an organization can share its resources and data with more precision, and quickly expand or limit a user's access as he or she takes on different roles.

Perimeter security will continue to serve as the first line of defense, but blind reliance on it will increase the risk of network breach and data interception. ZTA can gap-fill for perimeter security by catching network breachers at the next layer of access, matching the identity of users and devices to authorizations associated with those users and devices to ensure access is granted appropriately and securely. ZTA shifts the emphasis from the perimeter of a network to the discrete applications and services within a network, building more specific access controls to those specific resources. This method of wrapping security around applications and services is known as "microsegmentation," and it allows for more targeted security and management of access beyond traditional perimeter security.

By decreasing reliance on security around the perimeter of a network, the security of the data packets going in and out of a ZTA system becomes even more critical. As a result, encryption of those data packets is one of the most fundamental requirements of any ZTA, and should be standard in any good security architecture and data transfer (including protocols like SSL and TLS). The authentication and authorization process is also critical to the success of ZTA because it provides explicit verification of users and devices for each application and service, rather than implicit verification that assumes safety and grants access once a user and device has entered into the network.

Steps of Zero Trust

ZTA requires three fundamental steps, applied at the level of applications and services within the network:

1. Verify the user (authentication, part 1)
2. Verify the device (authentication, part 2)
3. Verify access privileges (authorization)

These three layers of verification are accomplished through a series of compliance checks based on the characteristics of each. These compliance checks can include information ranging from device encryption to user patterns of behavior, and can continue to expand as more information about users and devices is collected. Failing any of these checks will label the user or device as “non-compliant” and deny access to that user or device. While it is optimal to maintain a common set of compliance checks across a network and organization using modernized software systems, it is still possible to implement this critical piece of zero trust into networks that are dispersed and use legacy hardware and software, as in the case of DoD.

DoD Service-or agency-specific applications and systems can be built to have their own compliance checks, driving all users and devices that interface with them to adhere to those rulesets. This in turn can influence users and devices outside the orbit of those Services/agencies, which must similarly comply with those requirements to gain access. In this way, a ZTA can act as a virus (albeit a secure, trustworthy virus), “infecting” third party users or devices by forcing them to meet compliance standards if they hope to gain access. This in turn will make it easier for those third parties to eventually implement zero trust into their own applications and services, as they will already have certain standards of compliance on which they can build.

Compliance metrics can (and should) be iterative. ZTA can begin with a set of modest identity and device health checks, then develop a more sophisticated set over time depending on the sensitivity of the applications and services in question and the desired restrictions on access. Each “generation” of compliance checks will have the ability to add more nuance to access control as more information about users and devices is collected, building a more comprehensive digital picture of them. These checks will ensure not only the identity, but the integrity of users and devices: while identity checks will map the characteristics of each, integrity checks can run health diagnostics on devices to look for compromise. If a device has the

appropriate characteristics but demonstrates compromise (e.g., lack of encryption, virus infection), then the device will be deemed non-compliant and will be denied access.

One such identity check includes the management of certificates, which will help verify identity in ZTA by pairing public keys with the devices requesting access to an application or service. While certificates on their own may be stolen or imitated, ZTA protects against this risk by including certificates and public keys as one attribute for users and devices that must be corroborated with multiple other attributes in order to authenticate and grant a request for access. Additionally, certificates can be strengthened by including both user and device data into the certificate, requiring verification of both to enable access for either.

In advanced ZTA, this health inspection may ultimately be applied to the application or service being accessed - once the access requester has been verified, the zero trust model might then query the access destination to ensure the security of that destination. By transitioning from a “one-way handshake” to a “two-way handshake” model, an organization can better ensure that neither the users and devices nor the applications and services are compromised and infect one another.

More advanced forms of ZTA can also protect applications and services by applying microsegmentation to discrete areas within each application and service, building specific access rules for each of those areas. This methodology would protect against adversaries that penetrated the network perimeter and the application or service itself, limiting the data access and maneuverability of those adversaries. As in the case of microsegmentation at the application and service level, microsegmentation within those entities should require authentication and authorization specific to the subsegment in question, enabling a more precise “least-privilege access” model of security.

DoD Zero Trust Implementation:

DoD is Ready for Implementation

Many DoD network systems are old and siloed, but as noted above, that does not mean that they are limited from implementing zero trust. Zero trust solutions can start within a single organization or cross-organizational application, and rapidly drive all users and devices that interface with that organization or application to come into compliance and register their attributes for authentication and authorization.

ZTA can be implemented across DoD by building on an initial set of applications and services that adhere to zero trust principles, creating an ever-expanding web of users and devices that have their characteristics mapped for access management. Both organization-specific applications and cross-organizational applications are valid starting points for zero trust - as long as they have a broad number of touchpoints with users and devices, the basics of zero trust will spread across DoD.

DoD operations increasingly require better information flow between users and devices in and out of DoD, and ZTA can provide a solution for that need. DoD employees, contractors, third party subject matter experts and allied forces could more easily access and share information

using ZTA authentication and authorization rules to grant highly specific access without risking exposure of the rest of the network. DoD may already have the groundwork for an identification system to support the authentication process, in the form of the CAC card. The basic characteristics housed on a CAC can serve as a foundation for compliance, with the option of adding more nuanced characteristics over time. By leveraging existing identity management mechanisms like the CAC to shift to ZTA, DoD will enable improved data sharing and increase situational awareness while better defending against threats.

Due to the large number of siloed networks across DoD, any shift to ZTA would likely have to be incremental, starting with a standard set of identity checks for applications and services that could gradually be integrated into common mechanisms for authentication and authorization across DoD. As part of this effort, DoD will need to improve its digital management and tracking of user roles (and changes to those roles) across the organization in order to build access control for specific applications and services. While some of this effort will require security architecture reconfiguration, there will also need to be a shift in the security culture throughout DoD to promote accurate and consistent record-keeping of roles and other identity characteristics.

Questions to Ask When Implementing Zero Trust

Implementing ZTA involves a gradual shift across the entire network, services, and computing devices space. Implementation is a journey that can result in a successively more secure organization with each additional step taken. Below are a set of questions you can ask to see if your organization is effectively implementing ZTA.

1. In a world where the network boundary cannot be trusted, encryption of data in transit and at rest is critical. ***Does your organization have a plan to ensure that 100% of data transmitted between devices and stored on mass storage is securely encrypted? Does the organization have a robust and secure encryption key management strategy?***
2. Many recent cybersecurity exploits leverage the vulnerability of the network boundary to access incompletely secured systems accessible in the network. It is critical to ensure that piercing the network perimeter does not enable network access. ***Is your organization aggressively searching for and reducing traditional inside-the-network vulnerabilities, such as data stored in unsecured locations, excessive use of service accounts, and standing administration privileges?***
3. ZTA requires verification of the true identity of the user and device accessing a system. ***Is your organization enforcing Multi-factor Authentication on 100% of the services and servers available to users on all end user, devices including mobile devices?***
4. Strong ZTA requires verification that only healthy devices can access an organization's computing resources. ***Does your organization have processes in place to continuously scan all end user devices for malicious software, and is device health a real time criterion for accessing your organization's services?***

5. In ZTA, a user's identity alone should not provide access to the organization's services. These access rights are strictly applied depending on the user's role in the organization, and those rights should change whenever that user's role changes. ***Does your organization have a strategy in place to apply access rights strictly based on a user's role and to dynamically change those rights when a user's role in the organization changes?***
6. Ultimately, ZTA espouses that the network boundary provides no protection, and that the most secure services are those resilient enough to connect directly to the internet. ***Is your organization planning the migration of all critical services to directly connect to the internet and building the necessary security to enable these services to be secure in this internet-facing construct?***
7. As ZTA principles and design structures are introduced into a network, they should be built while bearing in mind the future scope and required maintenance of the network. ***Is your organization designing ZTA for scalability, supportability, and life-cycle management?***
8. While some DoD network requirements are unique to the Department, many others have solutions readily available in the commercial sector. ***Does your organization have a strategy for determining what you buy (commercially) and what you custom develop (or pay for custom development)?***

Where to Learn More:

Basics of Zero Trust:

- ACT-IAC, “Zero Trust Cybersecurity Current Trends,” April 2019, <https://www.actiac.org/zero-trust-cybersecurity-current-trends>.
- Arun Shresta, “A Guide to Effective SaaS Management Using a Zero Trust Security Model,” BeyondID, 2018, <http://info.bettercloud.com/n0KZ02k03wea0YZ0Sy101CB>.
- Doug Barth and Evan Gilman, “Zero Trust Networks,” July 2017, O’Reilly Media Inc., <https://www.oreilly.com/library/view/zero-trust-networks/9781491962183/ch01.html>.
- Google, “BeyondCorp Research Papers,” 2014-2018, <https://cloud.google.com/beyondcorp/#researchPapers>.
- John Kindervag, “Build Security Into Your Network’s DNA: The Zero Trust Network Architecture,” Forrester, November 2010, http://www.virtualstarmedia.com/downloads/Forrester_zero_trust_DNA.pdf.

Perimeter Security Vulnerability Case Study:

- Office of the Inspector General, “Cybersecurity Management and Oversight at the Jet Propulsion Laboratory,” June 2019, <https://oig.nasa.gov/docs/IG-19-022.pdf>.

Miscellaneous

- The National Institute of Standards and Technology (NIST) is currently developing a report on ZTA, slated to be published late 2019/early 2020. This report will provide a broader look at ZTA as it relates to networks across the federal government - once published, the NIST report may serve as useful additional reading to provide more context for this report.

Zero Trust Starter's Pack: Glossary

Authentication: “Authentication” is the identity confirmation for users and devices using attributes (e.g., name, location, time, etc.) universally recognized and tracked across the network. “Multi-factor authentication” (MFA) describes the process of cross-checking multiple attributes to verify identity.

Authorization: “Authorization” describes the access privileges granted to users and devices. Perimeter security authorization grants access to the network, while zero trust authorization grants access to individual applications and services within the network. Authorization is linked to authentication attributes, and can include “Role-Based Access Control” (RBAC) and “Attribute- Based Access Control” (ABAC), describe different approaches to authorization of user and device access. RBAC provides a common set of access rules based on pre-defined user roles, whereas ABAC provides access based on the combination of multiple user attributes which can include a user’s role.

Certificate authority: “Certificate authority” (CA) is a trusted entity that issues and manages digital certificates used for secure communication in a public network. CA confirms the ownership of public keys associated with the subject of the certificate. In zero trust, CA will serve a critical function in verifying content and content sources at the application and service level.

Identity and Access Management: “Identity and Access Management” (IAM) includes the processes, policies, and technologies that enable verification of users and devices requesting access to applications and services within a network. This includes the “authentication” and “authorization” mechanisms of ZTA. “Privileged Identity Management” (PIM) falls under the category of IAM and assigns varying degrees of access to different tiers of users, allowing the access controller to better monitor “Privileged Identities” (PIs) that have higher levels of permission.

Least-privilege access: “Least-privilege access” is a core principle of zero trust, and requires that users and devices only be granted access to the information, applications, and services absolutely necessary for them to function. Least-privilege access can also be referred to as “Principle of Least Privilege” (PoLP).

Key management: “Key management” refers to the protection, storage, back up, and organization of encryption keys used to protect data being transferred between end points. In the context of a zero trust network, encryption becomes increasingly important as a method of protecting data traveling across the network. Public Key Infrastructure (PKI) manages the creation and assignment of digital certificates and their accompanying encrypted public keys.

Microsegmentation: “Microsegmentation” is a method of creating secure zones within the network at the applications and services layer or below to allow better monitoring of lateral network traffic and management of access to those zones. Microsegmentation could enable zero trust principles by allowing administrators to control authentication and authorization at a more granular level than the network perimeter.

Perimeter model: “Perimeter model” refers to the traditional method of cybersecurity, which focuses on the defense of the edge of a network, or network “perimeter.” Once users and devices are granted access to the network, they are more easily able to move laterally into different applications and services within the network.

Secure Sockets Layer/Transport Layer Security Encryption: Secure Sockets Layer (SSL) and its follow-on, Transport Layer Security (TLS), are cryptographic protocols that protect data being transferred between end points by verifying connection security and performing “handshake” authentication between a client and a server.