

Risks to Zero Trust in a Federated Mission Partner Environment

Keith Strandell
Dr. Sudip Mittal

ABSTRACT

Recent cybersecurity events have prompted the federal government to begin investigating strategies to transition to Zero Trust Architectures (ZTA) for federal information systems. Within federated mission networks, ZTA provides means to minimize the potential for unauthorized release and disclosure of information outside bilateral and multilateral agreements. But when federating with mission partners, there are potential risks that may undermine the benefits of Zero Trust. This article explores risks associated with integrating multiple identity models and proposes two potential avenues to investigate mitigation of these risks.

INTRODUCTION & BACKGROUND

Within days following the cyberattack on the Colonial Pipeline, U.S. President Joseph R. Biden Jr., signed into effect Executive Order 14028: Improving the Nation's Cybersecurity.¹ Prompted by recent "sophisticated and malicious" cyberattacks, the order acts as a catalyst for federal agencies to take necessary and immediate steps to coordinate with industry on improving information sharing, adopting best practices, and migrating federal information systems from perimeter-based security to a Zero Trust Architecture (ZTA). The foundational elements of Zero Trust are micro-segmentation and a well-informed trust algorithm. When effectively implemented with data tagging, Zero Trust provides a strong compartmentalization model that lends itself to federated mission partner environments. However, in an environment where mission partners are responsible for bringing to the table their own identity models, consideration must be given to risks associated with federating multiple mission partners.

© 2023 Keith Strandell, Dr. Sudip Mittal



Keith Strandell is the Materiel Leader for the Royal Saudi Air Force AWACS Modernization Program and previously served as the Materiel Leader for life cycle management of enterprise applications within the Enterprise Services portfolio for the United States Air Force's Enterprise Information Technology (EIT) construct, including capabilities such as the Air Force's Office 365 instance in the Impact Level 5 environment. Keith started his career as a Communication and Information Officer in the Air Force, leading the Network Security and Information Assurance offices at Travis Air Force Base. From there, he transitioned to a Technical Program Management role leading various hardware and software acquisition/development efforts related to battle control, intelligence, and Enterprise Information Technology systems. Included in this time frame were four years managing Foreign Military Sales cases across the Pacific, European, and Central Commands. Keith is currently pursuing a Ph.D. in Computer Science & Engineering at Mississippi State University.

In this article, we investigate the risks associated with a multi-partner environment built on ZTA that federates with each mission partner's identity model. For purposes of isolating the impact of federated identities, the operating assumption is that the environment has fully implemented micro-segmentation and data tagging such that the primary risks are associated with the integration of multiple identity models. In addition to assessing the risks, we recommend two potential areas of investigation that may alleviate some of the risks associated with this architecture.

MISSION PARTNERS AND DATA PROTECTION

Combatant Commands (COCOMs) work with a variety of international mission partners, the most obvious being foreign militaries. However, there is a significant degree of cooperation that occurs with other agencies. In January 2010, U.S. Southern Command (USSOUTHCOM) responded to a request for earthquake relief support by Haiti. This Humanitarian Assistance and Disaster Response (HA/DR) operation required coordination with multiple international organizations, including foreign government agencies, nongovernment agencies, and foreign militaries. In order to share information effectively, data were kept unclassified to the maximum extent possible and public platforms were used for dissemination.² Another example of cooperation with international partners can be found in a recent partnering among U.S. Africa Command (USAFRICOM), the International Criminal Police Organization (INTERPOL), and local law enforcement from several West African nations. The operation targeted illegal fishing and "other maritime crimes" along the West African coast.³ Not only do these mission sets require sharing of unclassified data, but they also demonstrate the potential for both persistent and transient user bases operating in the same environment.



Sudip Mittal is an Assistant Professor in the Department of Computer Science & Engineering at Mississippi State University. He graduated with a Ph.D. in Computer Science from the University of Maryland Baltimore County in 2019. His primary research interests are cybersecurity and artificial intelligence. Mittal's goal is to develop the next generation of cyber defense systems that help protect various organizations and people. At Mississippi State, he leads the Secure and Trustworthy Cyberspace (SECRETS) Lab and has published over 70 journal articles and conference papers in leading cybersecurity and AI venues. Mittal has received funding from the NSF, USAF, USACE, and other agencies within the U.S. Government. He also serves as a Program Committee member or Program Chair of leading AI and cybersecurity conferences and workshops. Mittal's work has been cited in the Los Angeles Times, Business Insider, WIRED, the Cyberwire, and other venues. He is a member of the ACM and IEEE.

Attempting to create a collaborative environment to facilitate data sharing that allows for multiple missions and user bases increases the need for effective controls to prevent the unauthorized release and disclosure of information such as Controlled Unclassified Information (CUI). For example, data controlled as Not Releasable to Foreign Nationals (NOFORN) are not releasable to foreign mission partners; however, these data may need to reside in this environment due to a need to release to non-foreign entities such as the Federal Emergency Management Agency. Similarly, data controlled as "CUI//REL TO USA, FVEY" are releasable to members of the Five Eyes alliance.⁴ A comparable protection requirement exists for mission partner data. The Mission Partner Environment framework is designed to facilitate collaboration and sharing with "participants within a specific partnership or coalition."⁵ The implication is a requirement to ensure data are shared only within designated groups. For example, assume there are existing agreements among the United States, country A, and country B, as depicted in Figure 1. In this image, the overlapping areas represent shared data based on these partnerships. Each country contributing data to the environment expects the information it uploads to the system to be protected accordingly. That is to say, data transferred to the United States as part of a bilateral agreement with country A must not be released to country B without the express consent of country A.




Figure 1: Multi-country data sharing partnerships.

ZERO TRUST AND FEDERATION

The operating assumption in ZTA is that the network is compromised and therefore steps must be taken to minimize the potential impact of unauthorized access. Through micro-segmentation and data tagging, a ZTA can provide a framework in which compartmentalization is baked into the security model. The result is smaller trust zones, which reduce the potential for lateral movement of an adversary exploiting a vulnerability (see Figure 2). However, to realize the benefits fully, the ZTA must also implement a trust algorithm that takes in relevant data feeds to provide continuous authentication and authorization decisions on access requests. A robust trust algorithm will have access to contextual information on the requesting entity and device, the target resource, resource access policies, and threat intelligence.⁶ Access to these information feeds provides a more complete view of the request and associated risks. For example, consider the ability to access data related to the requesting entity’s device configuration to compare those data with data on known configurations and thus to predict the level of vulnerability associated with the device.⁷ Such an assessment increases the insight into the risk of a given request.




Figure 2: Lateral Movement in Perimeter Network vs. Zero Trust.

Federation in ZTA poses an interesting conundrum because, in an architecture that strives to remove trust, it introduces an inherent trust among the federated organizations. Here, the mission partners act as identity providers and are responsible for authenticating their users. Once authenticated, the identity is securely transferred to support the trust algorithm making the authorization decision. This model eliminates the need for the user to maintain security information related to a separate identity, which can reduce the risk of compromise associated with user behavior. However, it can undermine the rigor of the trust algorithm by preventing access to contextual information related to the requesting entity.

RISKS FEDERATING WITH MISSION PARTNER IDENTITY SOLUTIONS

Zero Trust touts a robust trust algorithm rooted in the ability to verify a user’s identity. However, in a federated model, the algorithm is only as strong as the weakest identity solution. Figure 3 depicts a simplified model of a federated ZTA environment. The network supports countries A, B, and C. The entities in each country (e.g., military, law enforcement) have their own distinct identity models that are federated with the system. The Registered

User List provides a means for restricting users, standardizing attributes, and providing redirects to the appropriate identity system for authentication.

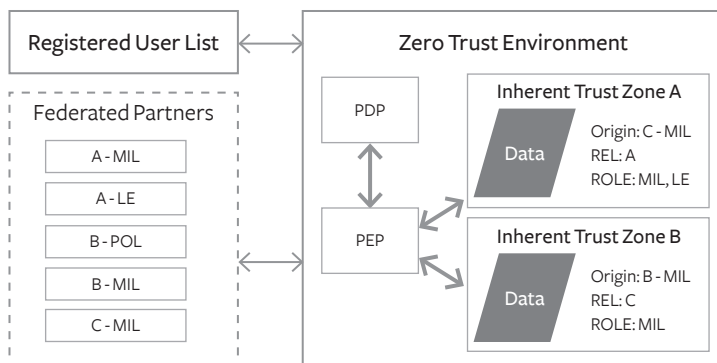


Figure 3: Notional Federated ZTA Environment

Ideally, each identity system would adhere to a minimum baseline that supports a context-rich authentication model. However, when balancing risk and mission requirements, mission may take priority and drive risky behavior or decisions. As such, there is the potential for integration with substandard models, which increases the risk of unauthorized access to the system. Assuming an ideal implementation of ZTA in the model above, access to a given Inherent Trust Zone will be restricted to an appropriate user base, and users granted access to a given zone cannot move laterally. Hence, a user who accesses Zone A above will not have access to Zone B without having gone through a separate access request. Therefore, in the model provided, the “C - MIL” identity system offers the widest potential reach for a threat, because it is the only one whose user base has authorized access to both Inherent Trust Zones.

Successful implementation of Zero Trust is predicated on a robust Trust Algorithm with access to contextual information around a given access request. For example, relevant information for an access request could include multiple authentication factors, device registration check, and device health status. The ability to access the contextual information around the requesting entity is a challenge in federated models.⁸ This model assumes contextual checks occur within the authentication pipeline managed by the mission partner, and therefore the ZTA environment is effectively blind to the degree of rigor used to authenticate a user. If it is assumed the authentication model for the “C - MIL” identity system is strictly a username and password, it becomes a prime target for adversaries looking to access the system. Given the strength of ZTA in containerizing information, an adversary should only have access to those Inherent Trust Zones to which the compromised account has access. In this model, the obvious impact of a compromised “C - MIL” user account would be the unauthorized disclosure of data in Inherent Trust Zones B and A. However, there is also the potential to upload misinformation and malicious code that could compromise entities in countries A and B.

The effective establishment and enforcement of a common, robust authentication process increase the security of the system; however, they do not address vulnerabilities in the supply chain. On December 8, 2020, it was discovered that SolarWinds had been compromised. The hack, attributed to Russia, affected approximately 17,000 SolarWinds clients, including several federal agencies such as the Department of Homeland Security and the Department of Defense. The attackers targeted a third-party vendor, Orion, that had a long-standing relationship with SolarWinds. Because Orion had been infiltrated, when clients of SolarWinds updated their software they inadvertently loaded malware onto their devices and thus gave hackers access to their networks, which in many instances resulted in significant data breaches.⁹ The attack is significant in that it focused on popular network infrastructure devices, which allowed for the vast attack surface. A comparable attack on the identity components used by “C - MIL” could provide an adversary with the ability to hijack existing credentials or bypass authentication processes. There also exists the potential for introducing fraudulent credentials, but, that can be mitigated with the effective implementation of a Registered User List.

This risk is amplified by strategic competitors’ ability to leverage the Diplomatic, Informational, Military, and Economic (DIME) framework to deliberately position state-sponsored technology that provides them covert access. Strategic competitors such as the People’s Republic of China (PRC) and Russia are actively exercising DIME strategies to advance their influence in regions around the world. General Stephen J. Townsend noted in his statement to the House Armed Services Committee that both countries have an “inside track” in central and southern Africa. He also stated that Russia is actively buying influence in the region and the PRC is investing billions in infrastructure and development in Africa.¹⁰ Within USSOUTHCOM’s area of responsibility (AOR), the PRC holds \$165B in loans and is using COVID-19 as a pretext to indebt nations in the region further while enhancing its integration with their infrastructure and technology. For example, as part of their COVID-19 response, the PRC was offering to donate Huawei technology.¹¹ The significant investments these competitors are infusing into the region provide the pretext to gain access to senior government officials with the leverage to secure deals that further embed their technology or allow insight/access to processes like identity management. The infrastructure investments in these regions serve, at a minimum, as a method to increase reliance and influence. However, they also introduce the potential supply risk noted above. Specific to the PRC, there are concerns related to the Military-Civil Fusion Strategy and how involved vendors such as Huawei are with the People’s Liberation Army and the extent of their collaborations.¹²

This concern is furthered by incidents which suggest not only security issues but the intentional inclusion of surveillance capabilities that lend themselves to espionage.¹³ While the Huawei push is focused on 5G, the concern extends to any presumably state-sponsored technology that may serve as critical infrastructure for mission partner networks. Coupled

with the potential for growing an insider threat, there is the potential to undermine the processes of the Registered User List and reintroduce the risk of fraudulent accounts.

POTENTIAL ENHANCEMENTS

Reducing the risks associated with federating multiple identity providers in the model depicted requires introducing an additional layer into the authentication process that provides contextual data. Two promising designs for consideration are blockchain and Adaptive Neuro-Fuzzy Inference System (ANFIS). The former shows promise in reducing the likelihood of compromised credentials while the latter has the potential to identify and flag behaviors that deviate from the norm.

Blockchain

Blockchain first gained popularity as the digital ledger supporting bitcoin transactions; more recent implementations have shown its promise as a mechanism for augmenting or replacing existing authentication systems. The strength of blockchain lies in its immutable, secure nature, which comes from the combination of Merkle Tree hashing, encryption, distributed architecture, and consensus protocol.¹⁴ Smart contract implementations of blockchain can support authentication models through its abilities both to store data and to automate processes. It has been proposed, for example, as an authentication model for a cloud-centric database that requires access from both internal and external users.¹⁵ Blockchain has been shown to be capable of storing digital identities and data necessary to support authentication. It has also been shown to be capable of authenticating devices in an Internet of Things (IoT),¹⁶ which may be leveraged to support an agent-based model that allows a user to register a limited number of devices. Within a federated network, blockchain has the potential to introduce a layer of managed context that decreases the likelihood of an account being compromised.

Adaptive Neuro-Fuzzy Inference System


Adaptive Neuro-Fuzzy Inference System is a machine learning framework that couples the learning capabilities of adaptive neural networks with the fuzzy inference system's ability to detect ambiguities in decision-making criteria. This combination makes it well-suited for applications such as nonlinear analysis, control systems, and expert systems. The ANFIS framework has been leveraged in areas such as an improving pattern password authentication performance for touchscreens,¹⁷ anomaly classification to support intrusion detection in a vehicular ad hoc network,¹⁸ and a continuous authentication system for mobile devices.¹⁹ The latter utilizes ANFIS to learn passive and active patterns of use for a given mobile user in order to define a behavioral model. This allows the authentication system to monitor behaviors continuously and support implicit authentication while also flagging deviations.²⁰

For the purposes of a federated ZTA environment, ANFIS has the potential to be leveraged in both a client-side and server-side model. A client-side model could potentially generate a confidence score specific to a user's behaviors that is passed as context for authentication. This could support identifying a compromised device. A server-side variant may support identification of anomalous behavior relative to an archetype based on attributes. For example, if a certain user base only logs in periodically to check email and a specific user's behavior is significantly more active, that anomalous behavior may represent an insider threat or compromised credentials.

CONCLUSION

The transition from perimeter-based cybersecurity to ZTAs should result in significant improvements in the overall security posture of enterprise networks. Specifically, it shows promise in the realm of multinational operations in which cooperation can often be born out of necessity and built on a tentative trust among mission partners. The inherent compartmentalization of a robust ZTA lends itself well to an environment rooted in mission partners' trust that their data are protected from unauthorized release and disclosure. Unfortunately, the benefits of Zero Trust can be undermined by the federating of multiple identity models, a risk made worse by actions of strategic competitors to employ the DIME framework to enhance their regional footprints, advance their influence, and deploy state-sponsored technologies. These activities increase the opportunities for social engineering, political influence, and clandestine cyber operations. Some of the risks can be mitigated by limiting federation to mission partners with known, trusted architectures and limited ties to strategic competitors while offering to host all other partners. This model, however, has the potential to be compromised when mission requirements outweigh the cybersecurity risks. To secure the environment's security posture further, additional measures should be investigated.

Two promising options for enhancing the authentication model that could be investigated as augmenting technologies are blockchain and Adaptive Neuro-Fuzzy Inference Systems. Blockchain gained popularity as the digital ledger supporting bitcoin transactions. However, recent efforts go well beyond that, using blockchain for authentication as part of a self-sovereign identity model. In 2019, a group of credit unions piloted the use of blockchain and noted the improvement in the authentication model could reduce a credit union's annual fraud expenses by \$150K just by reducing the authentication risks tied to call centers.²¹ ANFIS is a machine learning model that integrates adaptive neural networks with a fuzzy inference system. In a study on its potential use to support "continuous implicit authentication" on mobile devices, ANFIS was used to learn user behaviors for supporting implicit user authentication and identification of both informed and uninformed adversary attacks. While the model showed a 5% increase in user recognition, the improvement in informed adversary attacks was negligible and it underperformed on identifying uninformed adversary

attacks.²² The ANFIS architecture does show promise for user authentication on mobile devices. However, if paired with an identity model, it may be used as part of an enterprise authentication solution that focuses on learning archetype behaviors to identify when a user's behavior deviates from the normal behaviors of users assigned to the same role, or from the user's own behavior pattern. 

DISCLAIMER

The views expressed in this work are those of the authors and do not reflect the official policy or position of the United States Military Academy, the Department of the Army, the Department of the Air Force, or the Department of Defense.