Carnegie Mellon University
Software Engineering Institute

Zero Trust: Risks and Research Opportunities

Geoffrey Sanders
Tim Morrow

March 2021

CERT Division

http://www.sei.cmu.edu

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

DM21-0178

Table of Contents

List of Figures

Executive Summary

Zero trust (ZT) is a cybersecurity paradigm that focuses on resource protection and the premise that trust is never granted implicitly but must be continually evaluated. It encompasses a collection of concepts and ideas designed to enforce least-privilege resource access to systems and services. To understand what areas of ZT research could be beneficial for the Software Engineering Institute (SEI) to pursue, we developed a notional U.S. government agency. We used this agency to develop an operational vignette to help understand the nodes and actors that would interact with the agency in a hybrid cloud ZT environment.

After completing the vignette, we developed three mission threads to help understand how the nodes and actors worked at an abstract level, including assumptions, configurations, FedRAMP security controls, quality attributes, and thread steps.

The resulting mission threads highlight risks and research areas to consider for ZT environments. Our mission thread development identified five risk areas that should be considered for research:

- non-existent ZT standards
- known attacks
- dynamic risk policies
- assets
- technical debt

We recommend exploring these risks in four future ZT research areas:

- architectural modeling and analysis
- technology research
- reference implementations
- situational awareness

1 Background

Zero trust (ZT) is a cybersecurity paradigm that focuses on resource protection and the premise that trust is never granted implicitly but must be continually evaluated [Rose 2020]. It encompasses a collection of concepts and ideas designed to enforce least-privilege resource access to systems and services.

A ZT architecture (ZTA) is an enterprise cybersecurity plan that implements ZT concepts through component relationships, workflow planning, and access policies. To realize a ZTA, seven basic tenets must be followed [Rose 2020]:

1. All data sources and computing services are considered resources.
2. All communication is secured regardless of network location.
3. Access to individual enterprise resources is granted on a per-session basis.
4. Access to resources is determined by dynamic policy—including the observable state of client identity, application/service, and the requesting asset—and may include other behavioral and environmental attributes.
5. The enterprise monitors and measures the integrity and security posture of all owned and associated assets.
6. All resource authentication and authorization are dynamic and strictly enforced before access is allowed.
7. The enterprise collects as much information as possible about the current state of assets, network infrastructure, and communications; the enterprise uses this information to improve its security posture.

1.1 A Notional U.S. Government Organization

The lack of standards and guidelines for ZT implementation introduces risk for organizations looking to implement this cybersecurity paradigm. Therefore, we proposed investigating how implementing ZT in a notional U.S. government organization could introduce risk. We did this by developing a vignette and three mission threads for ZTA context analysis of the notional Bygone Pension Agency (BPA).

The BPA insures the continuation and maintenance of private, defined benefit plans. Organizations commonly file distressed benefit plan terminations due to financial insolvency, and the BPA assumes plan control and maintenance while reducing benefit payouts. Several entities interact with the BPA, including agency staff and contractors, financial institutions, government organizations, insurance providers, employers and practitioners, and workers and retirees. The BPA provides services to these entities through an agency datacenter, a Trusted Internet Connection (TIC) provider, and a cloud service provider (CSP). Figure 1 illustrates an overview of entities that interact with BPA.

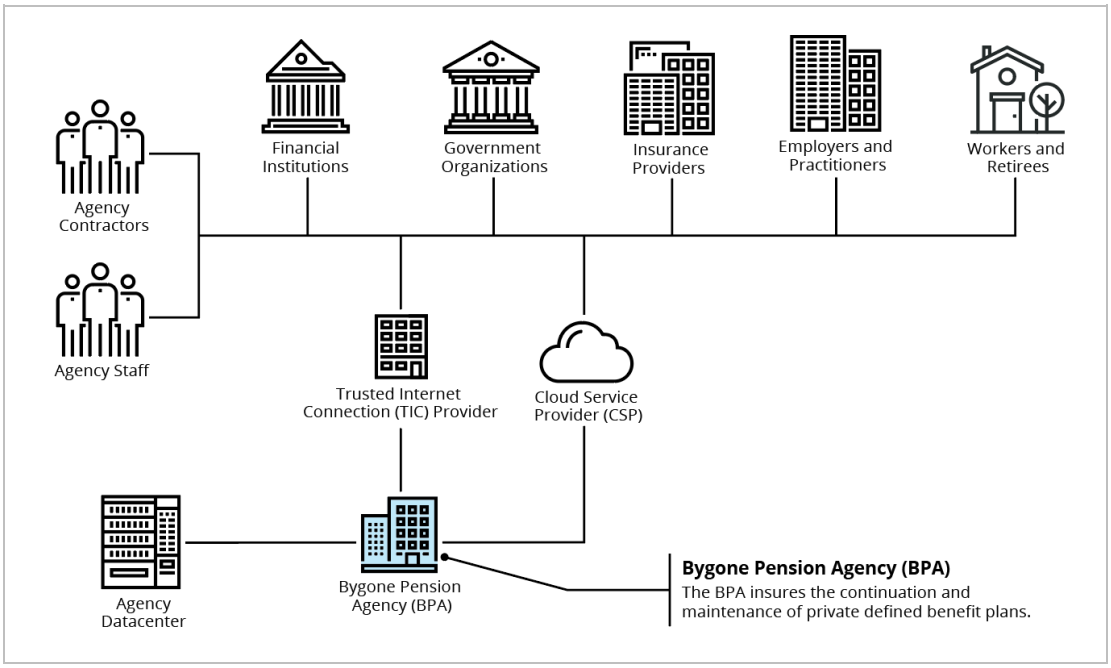

Figure 1: Bygone Pension Agency (BPA) Vignette

In this vignette, we assume that most organizations will transition from varied architectural patterns to ZT, and will not implement ZT tenets at the beginning of an architecture lifecycle. This approach implies that technical debt will determine if and when organizations can fully and appropriately adopt ZT tenets.

2 Risk Areas

BPA mission thread development and analysis identified several risk areas for organizations transitioning to ZT. These risk areas include system modeling, existing attacks, dynamic risk policies and assets. We assume that these risk areas are not exhaustive and merely provide a starting point for further research.

Because ZT is a set of concepts and not a defined set of standards, organizations implement unique organizational ZT solutions with common technologies.[1] When common technologies do not meet the requirements documented in ZT tenets, interested parties typically form groups to identify and develop solutions to meet them. For example, SPIFFE (Secure Production Identity Framework For Everyone) and SPIRE (SPIFFE Runtime Environment) are outcomes of the need to develop new standards and technologies to realize ZT tenets.[2]

In this section, we explore the five risk areas that we believe should be considered for further research.

2.1 Non-Existent ZT Standards

ZT is mainly a set of concepts; it is not a set of technology standards or products that can be architected and implemented. This gap introduces risk for organizations because they are forced to interpret ZT tenets. This situation is similar to a time before NIST 800-53 controls were created, and organizations were forced to interpret the meaning of *security*; everything was based on each organization's interpretation. The difference today is that U.S. government organizations are required to implement NIST 800-53 controls, while implementing conceptual tenets that are still not well defined.

Once controls are interpreted from ZT tenets, organizations will need to implement those controls with existing technology. Common technologies could fulfill ZT tenets,[3] or technologies could be developed specifically for ZT. This broad variety of technology leaves a large margin for error, which is likely to introduce vulnerabilities and interoperability problems. Problems could arise when one organization implements a tenet with technology extensions that are not recognized by another organization or product. For example, X.509 certificates are used in SPIRE for authentication and authorization. If an organization implements X.509 extensions that another organization or application does not recognize, incompatibilities between the two organizations could easily be introduced.

---

1 https://learning.oreilly.com/library/view/zero-trust-networks/9781491962183/ [Gilman 2017]

2 https://spiffe.io [SPIFFE 2021]

3 https://www.paloaltonetworks.com/cyberpedia/zero-trust-5-step-methodology [Palo Alto Networks 2021]

2.2 Known Attacks

Known attacks are a risk that may not be mitigated by transitioning to ZT. ZT effectively moves the organization's security perimeter from security zones (e.g., network boundaries or segments) to the service. This move transfers risk and security from perimeter controls to identity and access management (IAM) services because they are now responsible for making access and authorization decisions. However, even after this move, known web application attacks (e.g., session hijacking[4] and cross site request forgery[5]) could still be used in ZT environments.

For example, multifactor authentication (MFA) and federated identity are commonly used in ZT networks to implement IAM. Used together, they require multiple factors (e.g., a username, password, and token) for users to authenticate and access a service. Once access is granted to a user, they can execute workflows and functions associated with their identity. In web applications, attackers can "piggy back" on established web application sessions and execute commands that the user is unaware of but is authorized to do.

To illustrate this concept, imagine an IAM administrator authenticating to their IAM system to create user identities and access controls. If an attacker is monitoring this administrative function, they could create additional, unknown accounts using the administrator's IAM system connection. Similar risks occur with SSH multiplexing. Although SSH can be engineered for MFA, the session can also be hijacked once it is established.

To mitigate these risks, technology can be modified to limit the number of user sessions or connections to a service. However, these modifications might impact the user experience in ways that would result in removing these mitigations and negating the benefits of ZT adoption.

2.3 Dynamic Risk Policies

Dynamic risk policies are another ZT risk area. NIST SP 800-207 specifies that ZT must explicitly authenticate and authorize all subjects, assets, and workflows that make up the enterprise. To accomplish this requirement, organizations must develop and maintain dynamic risk policies and ensure that they are enforced correctly and consistently. However, standards for dynamic risk policies do not currently exist. Therefore, implementations of these policies could be specific to the vendor or system used, and these implementations may not be interoperable in environments such as hybrid clouds. For example, implementing dynamic risk policy solutions in Microsoft Azure could impact authentication and authorization with services in alternate CSPs, such as Amazon Web Services (AWS) or Google Cloud Platform (GCP).

Manual, human analysis of dynamic risk policies is not sufficient for the vast number of rules and services in a ZT environment. Therefore, systems and algorithms are needed for rule analysis, deconfliction, deployment, and management. The activity these systems create and log also requires

---

[4] https://owasp.org/www-community/attacks/Session_hijacking_attack [OWASP 2021b]

[5] https://owasp.org/www-community/attacks/csrf [OWASP 2021a]

automated security analysis. Gartner[6] and NIST SP 800-207 point to security information and event management (SIEM) systems as a solution for this needed analysis. However, these systems alone do not currently address all threats well; it is unlikely that they are sufficient for the increasing technical complexity that ZT introduces.

Organizations have not been able to automate enough to eliminate the need for human analysts. Combining this fact with long-term projections of cybersecurity worker shortages,[7] it is unlikely that less manual analysis will be required even when automation is a core principle of ZT. There will likely be much more data to analyze, but it is less likely that there will be an automated product to do this analysis effectively.

A security data analytics platform (SDAP) will likely be required. An SDAP is the logical big-data evolution of SEIM, and it will likely be required to analyze the large volumes of event data generated in a ZT environment. Each SDAP is uniquely built for an enterprise and requires a multi-disciplinary team of analysts and engineers to implement and maintain it.

2.4 Assets

Assets are a fundamental element for dynamic risk policies and policy decisions in a ZT environment. Unfortunately, asset management is a growing problem for organizations moving to the cloud.[8] Historically, organizations have not managed assets well in on-premise environments either.[9] Therefore, assets and asset management is another risk area for ZT adoption. ZT relies on automation for success; therefore, asset management must be comprehensive, inexpensive, scalable, efficient, and manageable through automation. DevOps and cloud service providers offer an advantage in these areas, where on-premise solutions typically fall short.

Currently, organizations commonly manage assets using X.509 certificates. These certificates are combined with user identities to determine access and authorization. Unfortunately, many certificate infrastructures are still managed manually,[10] which leads to system downtime and rogue or unknown certificates. To address this issue, ZT transition requires automated certificate creation, deployment, management, and analysis.

---

6 https://www.gartner.com/document/3912802?ref=solrAll&refval=263474518 [Gartner 2019]

7 https://www.herjavecgroup.com/wp-content/uploads/2019/10/HG-CV-2019-Cybersecurity-Jobs-Report.pdf [Herjavec 2020]

8 https://www.meritalk.com/articles/cybersecurity-and-asset-management-the-what-why-and-how-for-public-sector [MeriTalk 2020]

9 https://www.gartner.com/document/code/433878?ref=ddisp&refval=433878 [Gartner 2020]

10 https://www.gartner.com/document/3872966 [Gartner 2018]

However, one challenge that isn't addressed with this approach is unmanaged devices. Organizations that provide services to the general public are likely unable to easily install certificates to authorize and authenticate assets. This inability implies that alternative approaches for unmanaged assets must be identified and developed for ZT environments.

2.5 Technical Debt

Technical debt impacts an organization's ability to shift to ZT and introduces another area of risk. Applications and services might not have ZT features, and they might not be able to adopt them without significant additional resources. Since ZT is a concept and not a standard, a defined strategy for how to address this problem has not been developed. Therefore, general methods[11] should be used to address technical debt until ZT-specific approaches are needed or developed.

Potential solutions and tradeoffs will also be needed for risk identification, so larger groups can initiate risk management. For example, organizations may need to adopt Agile methods to update legacy applications with federated identity features. Web proxies may be able to help with implementing these updates, but non-web-enabled protocols would likely require significant modification.

[11] https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=546594 [Kruchten 2019]

3 Recommendations

We recommend focusing ZT research in the following areas:

- architectural modeling and analysis
- technology research
- reference implementations
- situational awareness

3.1 Architectural Modeling and Analysis

The risks we identified in Section 2 point to ZT architectural modeling and analysis as one area of research. Mission thread development and model-based systems engineering (MBSE) are structured methods that identify U.S. government risks through the application of ZT tenets. We believe hybrid clouds should be a specific architecture of focus for this work. Their complexity and need for interoperability suggest that they pose increased risk when compared to standard organizational networks interfacing with a single CSP.

One approach would be to identify an organization transitioning services to a CSP. This provides the environment to continue mission thread development and architectural models to identify gaps for further research. For example, our mission thread work in this scouting project identified that CSPs use agents to connect on-premise systems to cloud environments to manage both from a single interface. These agents do not appear to incorporate ZT tenets because they are implicitly trusted and do not have mechanisms that continually assess risks or threats. Deeper research into how to incorporate ZT tenets with the agents, along with threat modeling, could identify vulnerabilities, algorithms, or new methods that do not rely on X.509 certificates for asset management.

3.2 Technology Research

ZT technology is another area of research that we should explore. Reviewing technologies that organizations are using to implement ZT in hybrid cloud environments will help the SEI build a reference list for more extensive research. For example, NIST National Cybersecurity Center of Excellence (NCCoE) is using the completion of SP 800-207[12] to build a reference ZT architecture. Using mission threads from our previous example, we could identify how to implement this reference model with specific CSPs, in addition to risks that arise from them.

Another area of technology research that remains unanswered is how telemetry would be practically shared between U.S. departments and agencies (D/As) and the Cybersecurity and Infrastructure Security Agency (CISA). CISA TICs 3.0 artifacts call for sharing CSP telemetry, but details

[12] https://www.nccoe.nist.gov/projects/building-blocks/zero-trust-architecture [NCCoE 2021]

and standards do not exist. The SEI could research this area and identify architectures and services from CSPs that support telemetry sharing, in addition to their risks and threats.

A third technology research area would be to investigate how large organizations with large cloud infrastructures are approaching automated cybersecurity analysis and its applicability to ZT environments. For example, Capital One began the "Purple Rain" project,[13] which implements various commercial and open source projects to form an SDAP. Multiple technologies—including Apache Kafka, Apache Storm, Apache Metron, ElasticSearch, and ElastAlert—were integrated to build an SDAP for large-scale cybersecurity analysis. These technologies could be researched for automating ZT telemetry analysis.

3.3 Reference Implementations

A third area of ZT research to explore are ZT reference implementations. Results of architecture modeling, analysis, and technology research could be used to build a reference ZT implementation. This implementation could be done in collaboration with a U.S. government organization or within the SEI itself. The lessons learned from this research could be documented to understand the technology, requirements, and processes necessary for ZT transition and implementation.

Returning to the SDAP example from the Section 2, an SDAP reference implementation could be assembled to understand (1) what is required to implement one, (2) if there are general patterns that can be used across organizations and (3) the unique elements that do not transfer. Successful results could be brought forward as recommendations for CISA to implement a shared telemetry SDAP.

3.4 Situational Awareness

Another area of research to explore is ZT cybersecurity situational awareness (SA). ZT hybrid clouds are unique environments for security monitoring and SA. Contractual Cybersecurity Maturity Model Certification (CMMC) requirements for SA are emerging, but a defined approach to achieve them does not exist. CSPs claim that cloud services[14] meet these requirements, but empirical evidence to support these claims isn't available. Additionally, CSPs reference graph application programming interfaces (APIs) for security monitoring, indicating that custom situational awareness solutions will be required.

---

13 https://www.datanami.com/2017/04/21/purple-rain-bolsters-security-intelligence-capital-one [Woodie 2017]

14 https://devblogs.microsoft.com/azuregov/cmmc-with-microsoft-azure-security-assessment-situational-awareness-8-of-10/ [Banasik 2020]

Situational awareness research could be done in two areas previously mentioned:

- architectural modeling and analysis
- technology research

In previous SA research, Endsley identifies SA as both a state of awareness and processes used to achieve that knowledge.[15] Therefore, both areas, as they apply to cybersecurity and ZT, could be researched. We are not aware of any current SA research that addresses these areas in ZT hybrid cloud environments.

[15] https://www.researchgate.net/publication/292771806_Situation_awareness_analysis_and_measurement_chapter_theoretical_underpinnings_of_situation_awareness [Endsley 2000]

References

[Banasik 2020]
Banasik, T. J. CMMC with Microsoft Azure: Security Assessment & Situational Awareness (8 of 10). May 21, 2020. https://devblogs.microsoft.com/azuregov/cmmc-with-microsoft-azure-security-assessment-situational-awareness-8-of-10/

[Endsley 2000]
Endsley, Mica R. Theoretical Underpinnings of Situational Awareness. In *Situational Awareness Analysis and Measurement*. Lawrence Erlbaum Associates. 2000. https://www.researchgate.net/publication/292771806_Situation_awareness_analysis_and_measurement_chapter_theoretical_underpinnings_of_situation_awareness

[Gartner 2018]
Mahdi, D. *Technology Insight for X.509 Certificate Management.* Gartner, Inc. April 2018. ID G00345533. https://www.gartner.com/document/3872966

[Gartner 2019]
Riely, S.; MacDonald, N., & Orans, L. *Market Guide for Zero Trust Network Access.* Gartner, Inc. April 2019. ID G00386774. https://www.gartner.com/document/3912802?ref=solrAll&refval=263474518

[Gartner 2020]
Mahdi, D. & Collinson, D. Technology Insight for X.509 Certificate Management. Gartner, Inc. November 2020. ID G00433878. https://www.gartner.com/document/code/433878?ref=ddisp&refval=433878

[Gilman 2017]
Gilman, Evan & Barth, Doug. *Zero Trust Networks*. O'Reilly Media, Inc. 2017. ISBN: 9781491962190. https://www.oreilly.com/library/view/zero-trust-networks/9781491962183/

[Herjavec Group 2020]
Herjavec Group. *The 2019/2020 Official Annual Cybersecurity Jobs Report: A Special Report from the Editors at Cybersecurity Ventures.* https://www.herjavecgroup.com/wp-content/uploads/2019/10/HG-CV-2019-Cybersecurity-Jobs-Report.pdf

[Kruchten 2019]
Kruchten, Phillipe; Nord, Robert; & Ozkaya, Ipek. *Managing Technical Debt: Reducing Friction in Software Development.* Addison-Wesley Professional. April 2019. ISBN: 13: 978-0-13-564593-2. https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=546594

[MeriTalk 2020]
MeriTalk. *Cybersecurity and Asset Management: The What, Why, and How for Public Sector*. MeriTalk. September 2020. https://www.meritalk.com/articles/cybersecurity-and-asset-management-the-what-why-and-how-for-public-sector/

[NCCoE 2021]
Zero Trust Architecture. National Cybersecurity Center of Excellence (NCCoE) Website. February 2021 [accessed]. https://www.nccoe.nist.gov/projects/building-blocks/zero-trust-architecture

[OWASP 2021a]
Cross Site Request Forgery (CRSF). The Open Web Application Security Project (OWASP) Website. February 2021 [accessed]. https://owasp.org/www-community/attacks/csrf

[OWASP 2021b]
Session Hijacking Attack. The Open Web Application Security Project (OWASP) Website. February 2021 [accessed]. https://owasp.org/www-community/attacks/Session_hijacking_attack

[Palo Alto Networks 2021]
Palo Alto Networks. *Implementing Zero Trust Using the Five-Step Methodology*. Palo Alto Networks. February 2021 [accessed]. https://www.paloaltonetworks.com/

[Rose 2020]
Rose, Scott; Borchert, Oliver; Mitchell, Stu; & Connelly, Sean. *Zero Trust Architecture*. NIST SP800-207. National Institute of Standards and Technology. 2020. https://csrc.nist.gov/publications/detail/sp/800-207/final

[SPIFFE 2021]
SPIFFE website. February 2021 [accessed]. https://spiffe.io/

[Woodie 2017]
Woodie, Alex. *How 'Purple Rain' Bolsters Security Intelligence for Capital One*. April 2017. https://www.datanami.com/2017/04/21/purple-rain-bolsters-security-intelligence-capital-one/