Zero Trust Architecture: Risk Discussion

By: Alan Levine and Brett Tucker

Implemented well, Zero Trust Architecture (ZTA) promises to mitigate cyber risk for organizations of all sizes, risk postures, and cybersecurity maturity states. However, ZTA development, deployment, and operation present challenges that may hinder full adoption and sustained effectiveness and create new risk. Cyber risk should be evaluated by organizations as they make their decision for or against ZTA. Then, as organizations work toward full ZTA adoption and deployment that meets the criteria of maturity for the CISA Zero Trust Maturity Model, they should be aware of the risk that may not be solved by incremental steps. Finally, organizations should be prepared to address residual risk that may not be solved by their ZTA deployment, as well as new risk that may develop as they operate it. Guidance is available to support an organization's choice of ZTA and control the risk that may result from that decision.

The CERT Division at Carnegie Mellon University's Software Engineering Institute (SEI) has provided guidance related to the full spectrum of cybersecurity and enterprise risk and, particularly, the risk posed by a ZTA implementation, as this blog post describes. The essential tenets of ZTA and its attendant technologies and policies are not new. ZTA is new because of its focused, holistic approach to cyber defense that addresses every element of cyber risk with a unified program; as reinforced by ZTA, that program must operate exactly and relentlessly. Some organizations struggle to transition their cybersecurity architectures to ZTA, resulting in unresolved cyber risk. Fortunately, there are several strategies recommended by CERT products, including the CERT Resilience Management Model (CERT RMM) and the OCTAVE FORTE risk management model, that can assist organizations in managing risk as they seek to build, deploy, and operate ZTA as the core to their cybersecurity strategy.

The National Institute of Standards and Technology Special Publication SP 800-207, Zero Trust Architecture (NIST SP 800-207) identifies seven tenets of a ZTA program. Other agencies within the USG have also developed advice for organizations seeking to measure their degree of successful ZTA implementation, such as the CISA Zero Trust Maturity Model. This guidance is inherently valuable. However, even for organizations whose deployments are guided by it, challenges can remain for comprehensively embracing all the necessary elements for successful, sustainable ZTA implementations, incorporating all the solution sets necessary to address all their defensible assets, including their legacy systems. Organizations that move too slowly to deploy ZTA may fail in their mission to address existing and emergent cyber risk. Organizations that move too quickly to deploy ZTA may create new technical debt and, thus, new cyber risk; this risk would be in addition to the risk posed by any existing, unresolved debt that may be neglected or obscured as ZTA becomes a preeminent organizational objective.

A ZTA project should begin with a solid asset management program and assessment that identifies, with clear-eyed acknowledgement, the full array of an organization's cyber risk and an organization's cybersecurity state of maturity. Effective asset management can be challenging, but it is critical to understanding the nature of valued resources as much as related security requirements. A risk management process should be applied to assist in understanding an organization's realistic risk-based valuation of ZTA, the priority of ZTA among other organizational initiatives, and an appreciation of the necessarily competing organizational demands for resources and effort. Organizations should evaluate

the cyber risk that may be mitigated by ZTA and measure it according to the organization's risk appetite and threshold. This evaluation is important for all cybersecurity improvements, but ZTA is different: Its deployment and operation will necessarily become a focal point of the organization's cybersecurity efforts, consuming significant expense and capital, distracting from other potentially valuable initiatives within the organization's portfolio, and establishing an operational program that will require intensive attention to sustain.

The CERT process model, OCTAVE FORTE, provides a comprehensive roadmap for Enterprise Risk Management (ERM) policy and practices to help cybersecurity teams make better connections between the cyber risk they know and the organization's enterprise risk register. Like the principles of successful project management -- well-defined and followed scope, schedule, and budget -- good ERM programs rely on the acknowledgement of applicable risk appetite statements, the establishment of workable governance structures, and the careful development and articulation of policy and procedure. FORTE provides useful models for addressing each of these critical elements.

Regulatory concerns or other context, such as threat environment or risk profile, may help to make the case for an organization' risk-based decision for or against ZTA. The decision for ZTA requires careful consideration because this will be a strategic mission, not a tactical one. The impacts to normal operations may be many. The landscape of an organization's cybersecurity program will morph and, as it morphs, even while significant risk is being addressed, new risk may arise ; it will likely come from the usual suspects: Financial, where the organization's calculation of return on investment compels less budget than ZTA should require; and, Operational, where the normal conduct of the organization is impacted by unintended consequences.

Legacy assets that provide continuing value to the organization may not readily coexist within a ZTA architecture. Examples of legacy systems exist especially in capital-intensive industries that rely on operational technology (OT) to satisfy their organizational mission. The electric power sector is one example: The upgrade or replacement of significant portions of process instrumentation and control systems in order to satisfy the requirements of ZTA may be cost prohibitive and organizationally impracticable. Legacy IT and OT infrastructures may need more than software updates to satisfy ZTA's strict requirements for asset hygiene and user authentication and authorization; they may need to be replaced. New hardware and software to replace legacy systems would likely translate into substantial investment, and these costs may be compounded by lost organizational productivity and efficiency. The justification for ZTA should recognize these hurdles and be prepared to address them . Failure to address them may result in the failure of ZTA; break any part of it, and ZTA might break. All hope should not be lost in this case. Organizations must seek a systematic and methodical approach to evolve their security stack with an onset of compensating controls and eventual migration to long-term solutions.

In addition to OT and legacy asset realities, organizations that deploy ZTA may face other hurdles, including impediments to seemingly easy changes to policy, identity access and management processes, and new visibility measures. Each step of a ZTA implementation requires careful planning and change management, because even incremental steps toward ZTA may negatively impact normal operations. ZTA deployments that rely on bundled solutions, as many do to facilitate interconnectivity and compatibility, may miss out on the extra capability that might be provided by best-of-breed solutions; next generation solutions that bundle firewall, proxy, intrusion detection/prevention, identity management, and malware defense capabilities may be inherently complicated to configure and

operate, and any mistake in configuration or operation can create new risk even while the solution set works to mitigate existing risk. Endpoint solutions that bundle intelligence, defense, detection, and response capabilities may create similarly new risk.

One principle of ZTA capability, network segmentation, by itself may be a bridge too far for some organizations to deploy and operate. Segmentation is not only a practice of secure networking; it is also an essential part of robust high value asset protection. Available guidance describes the concept of segregated networks and how to deploy and manage them, and the necessary technology is available to support that effort, but segmentation remains a difficult process, subject to configuration errors and policy failures. Too often, there are more exceptions without compensating controls than rules for those controls. Every exception represents risk. Publications like NIST SP 800-37 the Risk Management Framework provide organizations with a standardized process to characterize their assets, identify controls, assess residual risk, and take additional action to accommodate their risk appetite. Organizations should recognize the overlap between the NIST Risk Management Framework in SP 800-37, the control sets prescribed in SP 800-53, along with the reference ZTA in SP 800-207 and utilize a mapping of those documents to guide their ZTA implementation journey. Organizations seeking ZTA are wise to follow an established construct as represented by these publications, and embrace an architecture that moves them toward more secure internal networking and better data protection. The road to these goals should begin with High Value Asset assessments. We should focus our programmatic efforts on the assets we value.

Visibility is another tenet of ZTA capability that may challenge organizations. Older network infrastructure simply may not tolerate expansive log collection and transmission that subverts normal computing and communication operations. Even for organizations with robust networks, persistent monitoring and validation may increase latency across the network. Visibility may result in other organizational conflicts. For example, policy shifts toward "bring your own device" in some organizations may constrain visibility. Change management strategies should emphasize workforce education as much as the actual physical, technical, and administrative changes needed to deploy ZTA. New monitoring, by itself, can create the perception among the user base that they are being watched, lose their trust, and create new insider risk for the organization.

In some cases, generous budgets and risk-averse appetites may drive an organization to attempt full scale adoption of ZTA. However, ZTA is best approached incrementally, because the architecture affects nearly all elements of a cybersecurity portfolio: edge, gateway, and cloud, internal network operation and segmentation, client, identity, applications/workloads, data, and – pivotally – program governance. ZTA implementation may conflict with other organizational priorities and may impact normal operations. Prior to embarking on implementation, organizations should evaluate the full organizational impacts ZTA will likely have, assign risk to these impacts, and then address that risk. Organizations should review their cyber risk appetite to understand the level and degree of investment that is acceptable to the organization deploying ZTA, and the time it will take to deploy it.

Organizations should consider employing the CISA Zero Trust Maturity Model to assess their current ZTA capability and define their desired state, as well as the path to get there. An organization that does not understand its current cybersecurity maturity state may fail at ZTA adoption, as may one that has not defined its ideal maturity state.

Once embraced as the core of an organization's cybersecurity strategy, a ZTA implementation should be planned with care. Organizations should make measured, risk-based decisions that are linked to specific outcomes, and this initial publication should represent a call for organizations to tailor means to measure the effectiveness of their decisions in reducing risk exposure. Here too, the CISA Zero Trust Maturity Model may provide overarching direction. At the minimum, the organization should consider these questions:

- Given the context of the organization, what level of cybersecurity is the organization seeking?
  - Does the organization have a documented risk appetite statement that provides quantitative and qualitative direction and informs its ZTA decisions?
- How does ZTA implementation directly relate to a demonstrable decrease in risk?
- With strengths in some ZTA-related capabilities and weaknesses in others, does the organization leverage its risk appetite to prioritize these new investments?
- How should the organization measure its implementation success?
- How should the organization measure the ongoing effectiveness of its operation?

Chief Information Security Officers (CISOs), Chief Risk Officers, and other leadership should define the organization's desired level of risk appetite so that deployment of incremental ZTA initiatives is prioritized to maximize the organization's ability to get them done, to digest them under a structured governance model, and to operate them so that they support the organization's target state of cyber maturity. Once a commitment to ZTA is established, leadership should advocate for resources, set expectations, and monitor for efficacy of implementation. The CERT Goal Question Indicator Metric (GQIM) model applies a formal process to assist an organization in setting and evaluating relevant metrics.

Lastly, an organization should consider the implications of ZTA implementation across its supply chain. Supply Chain Risk Management (SCRM) represents significant challenges to ZTA success because some suppliers may be unable to adopt or comply with some or all ZTA tenets. As a result, organizations should analyze and quantify (where possible) the risk that may be acceptable if, due to supply chain resistance, their ZTA deployment is incomplete and, thus, at least partially ineffective. Supplier contract terms and service level agreements may need to be reviewed and amended to mandate ZTA practices among suppliers, but such mandates come with their own impacts, some unintended: Some suppliers may decline to participate, they may modulate their performance, or they may demand substantial new contractual conditions to comply. The cybersecurity risk posed generally by the supply chain is often difficult to identify, quantify, and mitigate; supply chain risk to a ZTA program can be debilitating.

In summary, ZTA implementations should be viewed not as a sprint but as a marathon, with many incremental steps and substantive challenges along the way. Each challenge may create, hide, or magnify cyber and enterprise risk. Organizations contemplating ZTA should engage their ERM programs early to understand whether the effort to develop, deploy, and operate ZTA may result in unintended consequences, including new organizational risk. The impacts of ZTA to an organization's policy, processes, and technologies should be carefully managed to mitigate operational disruptions. Various resources, among them OCTAVE FORTE, NIST 800-207, and the CISA Zero Trust Maturity Model, applied thoughtfully and in concert, may support an organization in establishing necessary governance, advocacy, and understanding of interdependencies that successful ZTA implementations demand.

Acknowledgements

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.

DM22-0141

Bibliography

Cybersecurity and Infrastructure Division (2021), "CISA Zero Trust Maturity Model", Department of Homeland Security, https://www.cisa.gov/sites/default/files/publications/CISA%20Zero%20Trust%20Maturity%20Model_Draft.pdf.

Sanders, Geoff (2021), "Zero Trust Adoption: Managing Risk with Cybersecurity Engineering and Adaptive Risk Assessment", Software Engineering Institute. https://insights.sei.cmu.edu/blog/zero-trust-adoption-managing-risk-with-cybersecurity-engineering-and-adaptive-risk-assessment/

Caralli, Allen, White (2016), "CERT Resilience Management Model (CERT-RMM)", Software Engineering Institute, https://resources.sei.cmu.edu/library/asset-view.cfm?assetID=514489.

Tucker, Brett, (2020), "Advancing Risk Management Capability Using the OCTAVE FORTE Process", Software Engineering Institute, https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=644636.