ACALVIO | Booz Allen

WHITE PAPER

Advancing Zero Trust Maturity Through Cyber Deception

Detect and Counter Stealthy Threats Faster with Accuracy

In today's cyber threat environment, federal agencies and global businesses have no sure way to prevent adversaries from compromising their networks. That's why organizations are moving to adopt a Zero Trust Architecture (ZTA) that delivers data-centric security. A Zero Trust strategy assumes a breach, verifies identities, and grants the least-privileged access. But not all ZTA implementations are created equal: Defenders need all the visibility they can get to stay ahead of advanced adversaries.

Imagine an adversary has breached your perimeter defenses, gained a foothold in your environment, overcome prevention-based security measures, and started using stolen credentials and novel offensive techniques to move through the network toward the data they plan to steal. In this scenario, your ability to protect your most critical assets depends on how fast you can detect and respond to the attack. Standard sensors aren't designed to meet the need for such quick action. Emerging threats, meanwhile, are constantly accelerating.

Advanced adversaries are evading detection with sly offensive techniques like identity-driven attacks and AI-fueled polymorphic malware. In addition, they are using execution speed to rapidly propagate and compromise critical systems before the defense has the time to confirm threats. What's more, insider threats pose a growing security risk, with malicious activity by insiders appearing legitimate to anomaly and log analytics-based detection mechanisms. All in all, security operations center (SOC) teams are suffering from alert deluge: This reduces SOC reaction time and makes it easier for threats to spread.

Organizations need to reduce blind spots in their detection capabilities to stop attacks faster. In the context of the Department of Defense's seven-pillar model for Zero Trust, they need robust capabilities for "visibility and analytics." This is the hallmark of mature Zero Trust implementations. A leading way to achieve these capabilities—which can help turn the tables on stealthy hackers aiming to navigate the network undetected—is to include deception technology in your ZTA.

Deception technology involves predicting adversarial actions based on knowledge of malicious tradecraft, setting relevant traps of all kinds for the adversary, and tracking interactions with these traps. This approach goes hand in hand with an assume-breach mentality.

Incorporating Deception Technology in ZTA

Deceptions can take many forms, including decoys, baits, honeytokens, lures to protect identities, endpoints, network, applications, and data. Imagine, for instance, the deployment of identity honeytokens that are deceptive accounts deployed in identity stores and on endpoints. Any use of those accounts notifies an operations center that security has been compromised.

Overlaying a fabric of carefully chosen types of deceptions across enterprise systems, deciding the deception density, placing the traps at strategic locations of interest to the adversary, blending the ploys with the standard assets, and designing the ruses to be attractive to adversaries can make this strategy particularly effective.

Such an approach can help SOC teams reduce reaction time for threat confirmation, expand threat detection coverage, and uncover latent threats during threat hunting. Deception technologies not only provide clear signals of activity for investigation and action but also reduce data-processing costs and singular reliance on analytics for true/positive alerts. The high-confidence alerts raised by deceptions enable rapid response actions with defined processes tied to the "automation and orchestration" ZTA pillar.

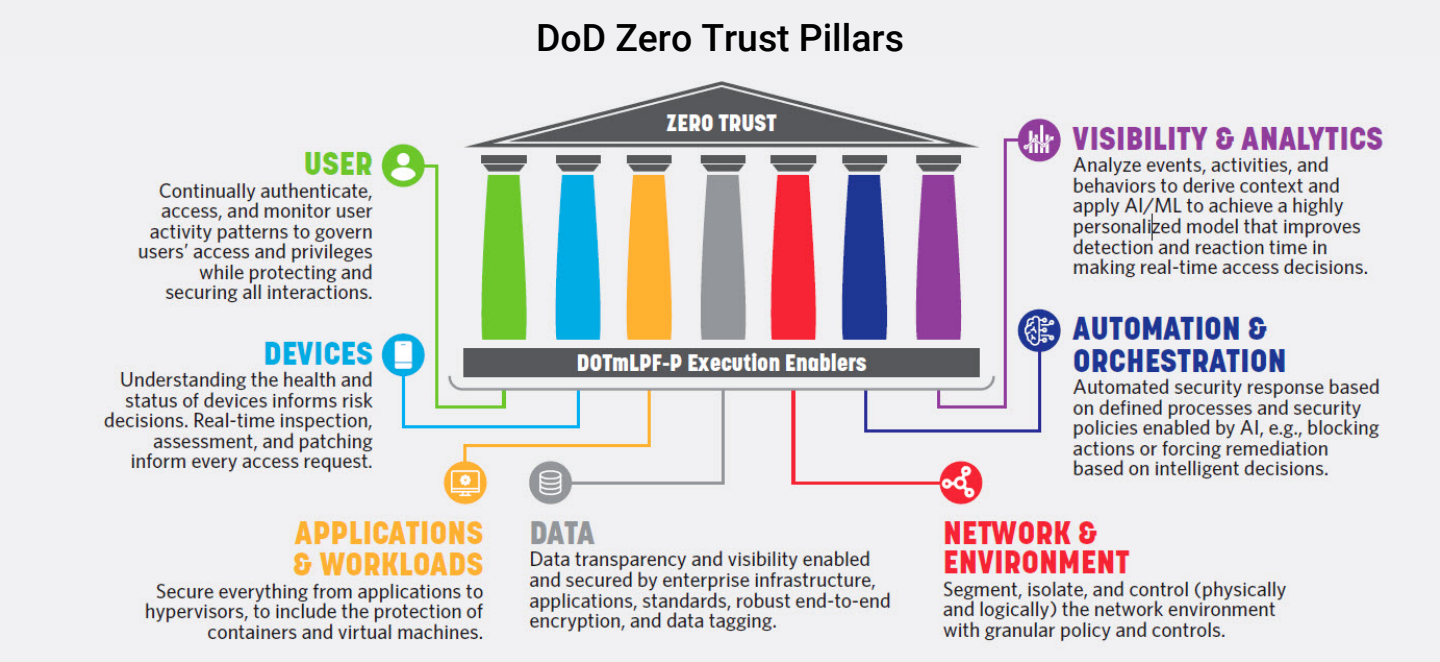

Source: DoD Zero Trust Strategy

AI/ML = Artificial Intelligence/Machine Learning
DOTmLPF-P = Doctrine, Organization, Training, material, Leadership and Education, Personnel, Facilities-Policy.

Use Cases for Deception-based Visibility in ZTA

Accelerate SOC Reaction Time

ZTA implementation increases alert volume for SOC teams and requires them to perform triaging of the alerts prior to threat investigation and response actions. Triaging involves manual alert correlation, deduplication and summarization. But there is a way to move beyond this cumbersome process to achieve better security outcomes.

Security teams can deploy, in strategic locations, a tailored set of deceptions (such as baits and honeytokens) designed to detect malicious activity in the early stages of attacks. This results in high confidence alerts that do not require triaging. With such early threat detection, the SOC can move more quickly to stop attacks from spreading and ensure protection of critical assets, and ZTA automation and orchestration capabilities can be brought to bear.

Expand Threat Detection

Adversaries use stealthy offensive techniques to bypass traditional sensors. After gaining a foothold in the environment, they perform offensive actions such as credential access, defense evasion, privilege escalation, and lateral movement. Traditional sensors require a certain level of signal to emerge before classifying the activity as malicious, giving the attacker an opportunity to propagate in the environment. Now, however, security teams can deploy deceptions to augment the traditional sensors and expand the detection coverage while also using the MITRE ATT&CK framework to characterize attack tools, techniques, and procedures (TTPs). Security teams can perform periodic coverage analysis to assess detection gaps based on traditional sensors.

The threats that are surfaced during the typical detection gap analysis include the following:

AI-fueled Attacks and Zero Days

Adversaries are using AI to generate polymorphic malware and ransomware variants. The polymorphic variants dynamically generate custom offensive code for each execution cycle, use memory execution techniques and avoid command and control (C2) communication to evade detection through traditional sensors. In addition, advanced persistent threats (APTs) use zero days to avoid detection.

A set of strategically chosen types of deceptions (decoys, baits, honeytokens) that are blended with the environment, made attractive for attackers to exploit and placed as an overlay on the critical assets, enables SOC teams to gain visibility into these evolving threats.

Use Cases for Deception-based Visibility in ZTA

Expand Threat Detection

Attacks on Identity Infrastructure

Mature Zero Trust strategies need more than prevention-based controls like identity access management (IAM) and multifactor authentication (MFA) to detect and respond to identity-driven threats. Adversaries obtain access to credentials from credential caches on endpoints, from identity stores, and from infrastructure servers that constitute the identity architecture. Adversaries use client-side attacks, offline attacks and identity zero days to stealthily compromise identities. Readily available penetration testing tools and offensive tools that target the identity architecture have lowered the bar for threat actors looking to perform identity-driven exploits. Adversaries use the credentials for lateral movement to gain trusted access to critical assets.

Traditional security sensors are unable to distinguish between legitimate and malicious usage of credentials. However, security teams can deploy identity honeytokens (deceptive user accounts, service accounts) in identity stores and on endpoints to detect identity-driven attacks.

Threats from Unmanaged Endpoints

Zero Trust environments are based on the principle of verifying the identity and granting least privileged access to resources from a device/endpoint. Securing the device/endpoint is an important step to ensure the right device gets access to the resources. But organizations are challenged to secure unmanaged endpoints—those endpoints with no endpoint detection and response (EDR) or extended detection and response (XDR) sensor deployed. Organizations have several types of endpoints that aren't compatible with EDR/XDR based sensors (e.g., legacy IT workstations or servers, bring your own device (BYOD) hardware, embedded devices such as printers and IoT cameras, and endpoints running custom OS versions).

Unmanaged endpoints are hard to patch due to limited access and legacy versions. This provides ready access for adversaries to compromise the endpoints through vulnerability exploits. Adversaries target unmanaged endpoints due to the lack of detection capabilities on the endpoint, providing opportunities to download and execute malware and spread across the environment.

SOC teams have limited and often no visibility to threats originating from unmanaged endpoints. To address this challenge, security teams can deploy deceptions to improve threat visibility from unmanaged endpoints as part of a mature Zero Trust implementation. Selecting the right type of deceptions—for instance, decoys impersonating unmanaged endpoints or baits on the endpoints—can help drive better security outcomes.

Use Cases for Deception-based Visibility in ZTA

Expand Threat Detection

Applications and Data-specific Threats

Adversaries target applications and data with vulnerability exploits to gain unauthorized access to sensitive data. Prevention-based controls are insufficient to protect applications and data, with over 20,000 unique vulnerabilities being reported every year, based on statistics from the vulnerability tracking database. Several of these are critical or high vulnerabilities. Patches often bring compatibility and stability challenges, creating long lead times for the patching process and creating a window of opportunity for attackers. As an example, the critical *Log4Shell* vulnerability that was discovered in December 2021 soon had a patch available, but organizations are still getting targeted by threats that successfully exploit unpatched systems.

Traditional sensors like EDR and XDR are endpoint-centric and are not aware of applications or data. Zero Trust is anchored around protecting sensitive data and critical applications. Organizations can deploy deceptions (decoy applications and data repositories), blend these with the environment and entice the adversary by projecting legacy or unpatched versions to create attractive targets in the eyes of adversaries. Baits and honeytokens can also be embedded as data deceptions in the standard data repository and applications. SOC teams gain visibility for threats targeting the applications and data as part of a mature ZTA implementation.

Threats that Bypass Traditional Sensors

Sophisticated adversaries are disabling traditional security sensors to evade detection, a tactic known in the MITRE ATT&CK framework as “defense evasion.” Adversaries obtain initial access to an endpoint, escalate privileges through a set of offensive techniques to gain administrative control over the endpoint. Adversaries use the administrative privilege to disable traditional sensors like EDR or XDR agents, deleting or reducing the logs and disabling other forms of agent-based monitoring. Ransomware and advanced malware threats are using defense evasion as an integral element of the attack lifecycle. SOC teams lose visibility to the endpoint.

Zero Trust strategies are based on continuous monitoring and visibility. Organizations can deploy deceptions on the endpoint that are data deceptions and are embedded in the standard data on the endpoint. The blended and passive data deceptions provide an independent detection control that is challenging for adversaries to find and evade. The added visibility strengthens overall Zero Trust maturity and enables SOC teams to detect and respond to threats more swiftly.

Use Cases for Deception-based Visibility in ZTA

Expand Threat Detection

Insider Threats

Insider threats come in different forms, including curious insiders and rogue administrators. Privileged insiders have trusted access to critical assets and sensitive data. Malicious activity performed by privileged insiders would not raise alarms in anomaly detection systems or result in suspicious log traces. This makes insider threats particularly difficult to detect using traditional security sensors.

Deploying deceptions such as data deceptions (baits) embedded in the data and on critical servers enables SOC teams to gain visibility to insider threats. This is a key capability for putting Zero Trust into action.

Threats to OT and IoT

Ransomware, insiders, and APT threats are targeting operational technology (OT) more and more. Adversaries know that once air-gapped OT systems are increasingly connected to the enterprise and beyond for governance, monitoring, administration, and security related reasons. What's more, they are finding various attack paths (e.g., they pivot from IT to OT environments or enter through internet-facing exploits).

Threat detection is particularly hard in OT environments. Traditional sensors provide SOC teams only limited visibility into OT threats. Endpoint tool coverage is limited by old equipment and vendor restrictions. The sensitivity of OT environments requires many tools to be passive, lest security sensors disrupt critical mission workloads. The embedded nature of OT assets and their specialized form factors introduce compatibility challenges for agent-based security sensors.

Applying Zero Trust principles to OT environments is a growing area of interest for the public and private sector. Now security teams can enhance visibility to OT threats by deploying deception technology to impersonate OT assets such as human-machine interfaces (HMIs), programmable logic controllers (PLCs) and controllers, as well as IT assets in the OT environment. The security team can safely introduce deception into OT environments to gain visibility into OT threats. The improved visibility for cyber threats better positions organizations to implement Zero Trust across OT environments

Use Cases for Deception-based Visibility in ZTA

Expand Threat Detection

Threats to Cloud Workloads

Misconfigurations are the most important threat vector for cloud workloads. Inconsistent usage of security settings, default configurations that are unchanged and configuration drift are common reasons for the misconfigurations. The adoption of multicloud workloads introduces more complexity, making it extremely difficult to secure cloud workloads based on prevention-based controls. Adversaries exploit the misconfigurations to gain access to cloud resources. Identity-driven attacks also are an important threat vector for cloud workloads, with adversaries targeting service accounts and secrets to gain access to cloud resources. Threat detection approaches in the cloud are limited mainly to scanning container images and cloud applications. Runtime threat detection options are limited and are specific for individual types of workloads in the cloud. Also, many types of cloud-native workloads—for instance, containers, serverless, and platform as a service (PaaS)—are not compatible with agent-based security solutions. The ephemeral and dynamic nature of cloud-native workloads poses an added challenge for threat detection and response, making it difficult to baseline and identify anomalous usage patterns.

Security teams can deploy deceptions (honeytokens, baits) that represent deceptive identities and data along with cloud decoys to detect cloud threats. Even security teams supporting organizations that are early in their Zero Trust journey for cloud environments can gain visibility to cloud threats to strengthen their overall detection capabilities.

Proactively Hunt Threats

Threat hunting is a proactive approach to identify threats in an organization. Threat hunting involves forming a hypothesis, searching for evidence based on the hypothesis and validating the threat to confirm the hypothesis. Traditional forms of threat hunting have relied on techniques such as indicator of compromise (IoC) sweeps and searches in logs. But advanced threats such as APT malware and stealthy ransomware variants lie dormant, waiting for a suitable opportunity to present itself. Uncovering such threats using traditional threat hunting techniques is challenging. Now, however, threat hunting teams can assess relevant types of deception options based on the hypothesis and deploy these in the target environment. As an example, the threat hunting team can deploy a decoy that projects a vulnerability and looks to the adversary like an opportunity for lateral movement. This offers a controlled opportunity for the latent threat to surface and enables hypothesis confirmation. Security teams gain improved visibility and threat intelligence. This, in turn, enables improved threat mitigation to support the Zero Trust strategy.

Next Steps for Improving Visibility in ZTA Implementation

Every organization has its own strengths and challenges when it comes to ZTA implementation. To position your organization's ZTA implementation efforts to achieve their full potential, leaders should consider three key questions:

1. **What are my organization's threat-detection gaps?** A red team assessment can uncover these gaps and provide a baseline understanding of current threat detection and incident response capabilities against real-world tactics, techniques, and procedures (TTPs) used by advanced persistent threats.

2. **Are there significant opportunities to rapidly close threat-detection gaps by taking greater advantage of deception technology?** A Zero Trust maturity assessment can reveal current strengths and gaps across the spectrum of ZTA capabilities. Some organizations now using deception technology may find new ways to scale up the effort and gain corresponding value. Other organizations may find incorporating deception technology is a completely untapped opportunity.

3. **What resources, plans, and other efforts are needed to rapidly improve visibility across the ZTA?** Assessment results can help leaders make informed decisions about investment priorities and other plans designed to close threat-detection gaps.

This content was created jointly with Booz Allen Hamilton and Acalvio Technologies. It was initially published at https://www.boozallen.com/expertise/cybersecurity/advancing-zero-trust-maturity-through-cyber-deception.html

Acalvio's mission is to help enterprises actively defend against advanced security threats with precision and speed. Acalvio's patented ShadowPlex Cyber Deception platform enables organizations to detect, engage, and respond to malicious activity across IT, OT, and cloud environments.

For more information, please visit www.acalvio.com

Trusted to transform missions with the power of tomorrow's technologies, Booz Allen Hamilton advances the nation's most critical civil, defense, and national security priorities.