Openpolicy

Sep 30, 2024

To: The National Institute of Standards and Technology (NIST), via nccoe-zta-project@list.nist.gov. *Atten: cherilyn.pascoe@nist.gov*

OpenPolicy's comment on
NIST Special Publication 1800-35- Implementing a Zero Trust Architecture

OpenPolicy appreciates the opportunity to provide feedback and suggestions to NIST on the SP 1800-35 Zero Trust Architecture (ZTA) guidelines. As a technology company dedicated to democratizing the engagement of innovative enterprises with policymakers, OpenPolicy is deeply committed to supporting the implementation of this crucial framework. We emphasize the importance of leveraging advanced technologies and innovative practices to enhance policymaking, streamline compliance, and fortify cybersecurity governance within the ZTA context.

We anticipate active engagement with NIST and related federal agencies as the implementation of the ZTA progresses and stand ready to offer additional insights as necessary. We are dedicated to supporting the cybersecurity objectives outlined in the ZTA, which are fundamental for guiding organizations in establishing secure digital infrastructures, managing cyber risks, and fostering a resilient security posture across AI, IoT, and OT ecosystems.

We believe that an open and collaborative policymaking process is essential for advancing the effective implementation of the ZTA. The involvement of innovative companies, particularly startups at the forefront of developing cutting-edge security solutions, is critical to this endeavor. These companies contribute technologies that align with the ZTA's requirements and relevant NIST guidelines, including frameworks for securing endpoint devices, implementing AI-driven analytics, and enhancing security across interconnected systems. OpenPolicy is committed to actively collaborating with these forward-thinking communities to ensure their expertise and innovations significantly contribute to the nation's cybersecurity initiatives. The technologies developed by these companies, including Armis, ADAMnetworks, Lasso, Kiteworks, HUMAN, Finite State, Cranium, InfoSec Global, Vaultree, and others, are instrumental in evaluating and securing essential cybersecurity components, from data protection and endpoint integrity to threat detection, response, and AI readiness within the ZTA framework.

To fully realize the potential of a secure and resilient digital infrastructure, **the ZTA should embrace the integration of AI, IoT, and OT-specific security controls within its policy**

engine's decision-making processes to proactively mitigate the evolving risks associated with AI-driven threats and IoT vulnerabilities. Simply establishing these controls is not enough; their real-time application and effectiveness across interconnected systems are paramount. By leveraging AI and IoT-specific Policy Information Points (PIPs), the ZTA guidelines can achieve continuous threat monitoring and response, effectively addressing the complexities and challenges posed by these rapidly evolving technologies.

Incorporating tailored endpoint security measures, comprehensive checks, and controls that address the unique vulnerabilities present in OT and IoT systems is essential for adapting the ZTA model to future threats, such as those posed by AI, PQC, and data misuse. Integrating **quantum-resilient cryptographic measures** into resource management, for example, is crucial to ensure that the architecture aligns with emerging quantum-readiness strategies and maintains preparedness against quantum threats.

As the ZTA guidelines mention, a unified approach to continuous access management, incorporating multi-factor authentication (MFA) and adaptive authentication tools, is necessary to ensure ongoing verification for all entities—human and machine—throughout their interaction with enterprise resources. This **approach reinforces a Deny by Default state, where only verified devices are allowed connections, aligning with core ZTA principles.** This level of identity protection is specifically critical for the AI, OT, and IoT ecosystem, where adaptive, real-time responses are required to counter evolving threats effectively. Furthermore, integrating advanced, AI-driven analytics within Policy Information Points (PIPs) is fundamental to detecting, assessing, and responding to threats in real-time across IT, OT, and IoT systems. This integration enhances ZTA's resilience against sophisticated attack vectors and leverages dynamic verification technologies and capabilities to preemptively block advanced threats, including direct-by-IP connections to Command & Control (C2) operators. By emphasizing AI-driven adaptive threat intelligence and ensuring that connections are strictly limited to verified, trusted sources, organizations can establish a more robust and true proactive security framework and approach within their ZTA implementation.

As the architecture evolves, secure data exchange across IT, OT, and IoT domains must be facilitated by **adopting encryption, tokenization, and access control measures, in alignment with NIST SP 800-213A guidelines,** to ensure the protection of sensitive data across diverse environments. To proactively protect and ensure the security of sensitive data as it moves beyond traditional network boundaries, it is essential to **adopt a content-defined ZT model**. This approach prioritizes securing the data itself by addressing privacy and compliance risks through a unified approach that consolidates and controls all sensitive content exchanges.

By implementing advanced data control measures, such as encryption and centralized visibility across email, file sharing, and other communication channels, organizations can enforce robust and consistent ZT security policies. This ensures that sensitive data is protected both at rest and in transit, while aligning with key regulatory frameworks like GDPR, HIPAA, and NIST 800-171. Moreover, a unified content management approach provides real-time monitoring, audit trails, and risk assessments, supporting continuous verification processes and enabling rapid responses to emerging threats, which enhances the adaptive ZT practices and mindset required for robust security.

Engaging in collaborative efforts with agencies such as CISA and OMB will be pivotal in advancing a consistent approach to cloud, data, and infrastructure security. By enhancing these aspects, the ZTA will be well-positioned to support a robust, AI-ready, and secure architecture that aligns with national cybersecurity frameworks and mandates.

We look forward to collaborating with NIST as the ZTA implementation advances and providing further expertise as needed. Our commitment lies in supporting the ZTA's cybersecurity goals, which are crucial for guiding organizations in building secure digital infrastructures, effectively managing cyber risks, and cultivating a resilient security posture across AI, IoT, IT, and OT ecosystems. We believe these objectives are fundamental to enhancing overall cybersecurity in our increasingly interconnected digital landscape.

1. **Enhancing Zero Trust Architecture with Advanced AI, IoT, and OT Security Solutions**

With the evolving ZTA threat and vulnerability landscape, it is essential to extend discovery and inventory processes beyond traditional IT assets to encompass AI systems, IoT devices, and OT components. These interconnected assets introduce unique vulnerabilities, which, if left unidentified, could be exploited by attackers. Utilizing AI-driven discovery tools ensures a comprehensive understanding of all assets, even those that are not easily detectable, thereby supporting a more robust ZTA. **Armis is a technology example with its advanced AI-driven discovery tools. It offers comprehensive asset identification capabilities that classify and map OT devices by function, location, and criticality, ensuring a thorough and accurate inventory and supporting a more robust ZTA implementation.**

Extending beyond traditional IT to cover AI, IoT, and OT environments ensures the enforcement of least privilege principles in formulating access policies. Armis's enhanced security features—such as tracking asset locations and custom tagging—enable organizations to extend ZT principles beyond traditional IT, covering AI, IoT, and OT environments, which allows the creation of detailed access policies tailored to dynamic, interconnected environments. Furthermore, as organizations assess existing security

capabilities, Armis's AI-based analytics and anomaly detection tools enhance threat detection and response. This ensures that ZTA components remain adaptable to sophisticated threats, aligning with NIST standards and FISMA compliance requirements.

A **risk-based approach to eliminate gaps in ZT policies** should prioritize protecting AI models, training data, and IoT/OT data flows, as these often contain high-value information vulnerable to adversaries. Implementing stringent access controls and segmentation around these assets would ensure effective mitigation of data exfiltration, tampering, or unauthorized access risks, a unique capability Armis offers by leveraging network analysis capabilities to identify IP addressing and protocols. Armis's Asset Intelligence Engine provides detailed device information, acting as a foundational element that supports the gradual rollout of ZTA and ensures security measures evolve with emerging threats.

Verification of ZTA should encompass real-time threat detection, compliance monitoring, and adaptive access policies for AI and IoT devices. It is crucial to leverage real-time alerts on vulnerabilities and risks to enhance continuous verification, supporting the maintenance of ZTA outcomes and ensuring rapid response to threats. **The continuous evolution and improvement of ZTA is a testament to the need for adaptability in the face of evolving attack patterns**. To that end, Armis's AI-driven monitoring tools adapt to evolving attack patterns. Its dynamic security information sharing and quantum-resilient encryption standards ensure that ZTA remains effective in responding to emerging threats, regulatory requirements, and technology advancements.

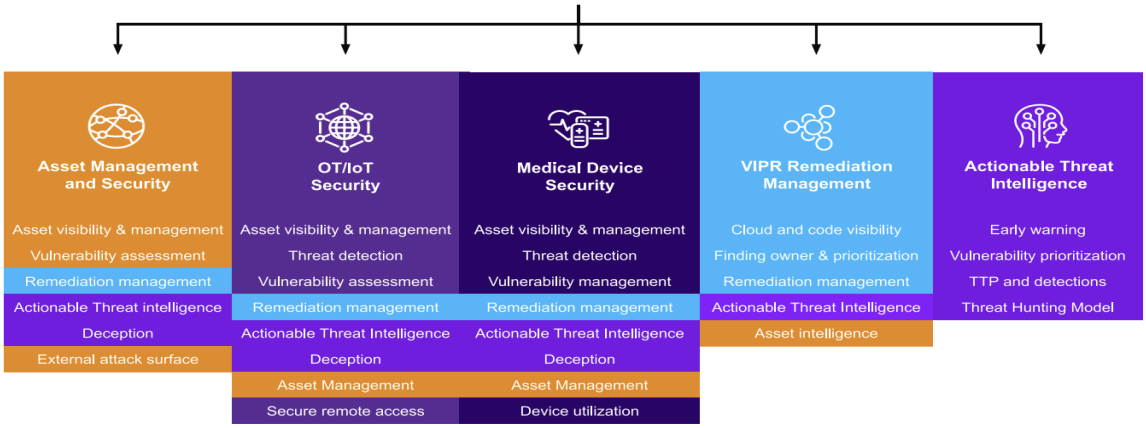

(Armis CentrixTM Suites empower organizations with a robust features capabilities and scalability to comprehensively address the entire attack surface, Source: Armis)

2. AI Risk Management and ZT Implementation

As emphasized throughout the ZTA document, it is essential to expand risk assessment frameworks to address not only technical and non-technical risks but also the unique challenges posed by AI systems, including data privacy, algorithmic bias, fairness, explainability, and the potential impacts on underserved communities. The guide, particularly subsection 8.4, underscores the necessity of integrating these considerations throughout the AI lifecycle, from training to application stages. **Cranium AI serves as a technology example with its advanced capabilities in AI risk management, offering comprehensive model protection controls that address vulnerabilities both during training and application phases.** This includes considering threats LLMs pose and mechanisms to control which data or unstructured content is introduced into these models.

An automated approach to risk reporting, testing, and mitigation significantly enhances the ability of organizations to establish **transparent communication of identified risks and mitigation strategies with stakeholders such as agency leadership, policy experts, and affected communities.** Cranium AI's platform, with its unique features, provides this capability, fostering trust and accountability in AI cross-sector participation. It ensures that measuring an organization's security posture is closely tied to accountability measures, allowing for the automation of software-based approaches to secure AI models and identify risks without disrupting existing workflows or data pipelines, enabling organizations to maintain robust AI model protection.

In alignment with subsection 8.1, it is crucial to emphasize the need for a **comprehensive inventory of AI assets, including datasets, models, experiments, and deployments.** This requirement is effectively addressed by Cranium AI's Enterprise Software platform, which is designed to maintain a thorough inventory that aligns with AI OMB requirements. The platform not only maps these assets into development pipelines but also offers continuous monitoring and alerting for key adversarial threats, such as data poisoning, model backdoors, evasion, and inversion. This approach ensures visibility, compliance, and transparency across AI environments while also supporting supply chain visibility and facilitating compliance reporting.

Cranium's platform offers robust capabilities to map, monitor, and manage AI/ML environments against adversarial threats, allowing organizations to gain visibility, security, and compliance across their AI systems without interrupting model workflows or data pipelines. This comprehensive approach enables organizations to gather and share information about the trustworthiness and **compliance of their AI models with third parties and regulators quickly and efficiently,** which further reinforces and implements ZTA guidelines and emphasis on risk strategies and third-party management within ZTA.

(AI/model risk measurement in support of compliance and control assessment; Source: Cranium AI)

3. AI-Driven and Adaptive Approach to ZTA Implementation

ADAMnetworks represents the next generation of ZTNA Zero Trust Connectivity (ZTc) solutions by leveraging AI-driven dynamic policies, adaptive threat intelligence, and sovereign data custody. This evolution aligns seamlessly with modern cybersecurity frameworks, effectively addressing the limitations of predecessor technologies while reinforcing robust, scalable, and real-time adaptive security measures for today's complex threat landscape.

Subsection 2.2 highlights the need for adaptable, real-time solutions that work across organizations of all sizes. ADAMnetworks' adam:ONE® addresses these challenges with its AI-driven dynamic allow-listing, providing **a scalable solution for both SMBs and enterprises.** Its **default-deny-all policy reduces the attack surface, offering protection before asset identification**, while full Layer 2 **visibility resolves foundational issues like asset management and visibility.** Administrators gain real-time access to logs, allowing for effective asset monitoring and control. Adam:ONE® supports multiple modes, from reactive threat intelligence to proactive ZTc mode. Its Smart rollout enables a gradual transition into ZTc, with AI-generated allow-lists ensuring minimal disruption. The solution's flexibility in policy enforcement at the group, network, or device level makes it adaptable to each organization's unique needs, **offering a tailored path to achieving ZTA maturity.**

Addressing section 5, General Findings, the regulation rightly emphasizes the need for a proactive security posture and continuous verification within a ZTA. However, it falls short in addressing how to maintain real-time, adaptive defense mechanisms against evolving

threats. ADAMnetworks exemplifies this proactive approach through its AI-powered dynamic allow-listing, which ensures a real-time, adaptive response to emerging threats, filling the critical gap in maintaining a ZT posture that evolves alongside threats. Unlike traditional Zero Trust Network Access (ZTNA) solutions that operate under an "Abdicated Data Custody" model—requiring a third-party man-in-the-middle to decrypt and inspect all traffic—ADAMnetworks adopts a **"Sovereign Data Custody" model**. This approach not only ensures that enterprises retain full ownership and control over their data but also requires only the source and destination information to enforce proactive protection, all without the need to break encryption, which provides a **crucial advantage in today's increasingly encrypted and complex network environments**, maintaining security without compromising data integrity.

As part of the guidelines's goals of establishing and maintaining a robust ZTA, ADAMnetworks's Don't Talk to Strangers (DTTS®) technology serves as a cornerstone in this effort, offering dynamic verification that actively thwarts advanced threats before they can infiltrate. By adopting a 'Deny by Default' approach, DTTS aligns with the regulation's emphasis on stringent access controls and real-time threat detection, which are essential in today's evolving threat landscape. This technology exemplifies **how dynamic verification can preemptively identify and block advanced threats, such as direct-by-IP connections to Command & Control (C2) operators—an area where traditional Security Operation Centers (SOCs) often encounter challenges**. The ability of DTTS® to enforce strict verification without interrupting business activities demonstrates how adaptive, real-time security measures can be practically integrated within ZTA frameworks, ultimately fostering a more adaptive and responsive security posture in line with regulatory expectations, which ensures that organizations do not just establish a ZT posture but maintain a resilient security framework that evolves alongside threat landscape.

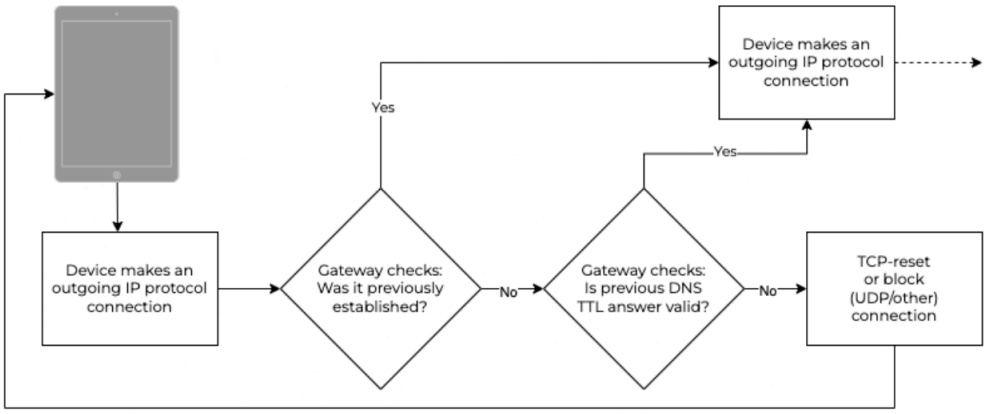

(Don't Talk To Strangers (DTTS)- A stranger in this context is an IP address that has not been resolved via a permitted DNS query- Source: ADAMnetworks)

Moreover, while the ZTA guidelines on multi-vendor integration lack specific guidance on achieving seamless interoperability, ADAMnetworks' integration-agnostic approach, powered by real-time threat intelligence aggregation through DNSharmony, effectively addresses this need. By utilizing a "muscle-brain" distributed protection model, ADAMnetworks combines the efficiency of on-premise gateway-hosted clients with centralized cloud control, **enabling organizations to implement a cohesive and dynamic ZT model across diverse technologies, ensuring seamless integration and strengthening their overall ZTA.**

While section 7 effectively emphasizes the need for comprehensive identity, authentication, and authorization measures in risk and compliance management, it must also address how to implement these measures in real-time or adapt them to an evolving threat landscape. Achieving true resilience and ensuring individualized resource protection requires real-time threat detection and mitigation, which aligns with **ZTA's core principle of minimizing the attack surface and protecting critical assets.** ADAMnetworks addresses this gap with its robust ZTc capabilities, delivering continuous risk assessments and real-time verification intelligence to ensure compliance with NIST SP 800-53 controls and EO 14028 security measures. Its DTTS® technology further enables immediate threat mitigation, effectively safeguarding resources and minimizing potential attack vectors, thereby enhancing the regulatory framework's goal of individualized protection within a ZTA.

Formulating access policies that adapt to mission-critical requirements, as outlined in subsection 8.2, underscores the necessity for these policies to dynamically adjust in response to evolving user behaviors and emerging threat intelligence. This demands an AI-driven approach capable of continuously adapting to shifting behaviors and threat landscapes, ensuring that access controls remain adaptive and resilient. ADAMnetworks' AI-driven Dynamic Allowlisting and DTTS® technology ensure that access controls remain both effective and responsive, adapting in real-time to changes in user behaviors and emerging threats. **This capability goes beyond traditional policy enforcement, aligning perfectly with principles like least privilege and separation of duties, and reinforces the need for dynamic access policies that evolve with organizational requirements** – an aspect that the regulation could emphasize more explicitly.

As organizations aim to leverage existing security capabilities, as detailed in section 8.3, seamless integration within a ZT framework is crucial. However, this section does not provide a roadmap on how to achieve this integration effectively. ADAMnetworks offers a practical, **cost-effective path by seamlessly integrating with existing infrastructure through adaptive deployment strategies,** enhancing ZTA capabilities without necessitating a complete overhaul. This approach ensures organizations can enhance their ZT capabilities while retaining their current investments, reflecting the evolution from

previous network protection methods that sought to centralize traffic redirection via proxies.

Regarding section 8.7, which emphasizes continuous improvement and evolution, the regulation's guidance on maintaining an up-to-date ZTA should be more specific on **proactive solutions.** Rather than relying on traditional integrations, such as identity-based systems, **ADAMnetworks' approach bypasses this need by automatically triggering a deny-by-default state**, treating each connection as potentially malicious. Leveraging AI, it dynamically verifies and determines what to allow through versus blocking potential threats. This proactive monitoring and advanced threat intelligence capability not only adapts to the evolving threat landscape but also provides robust protection against zero-day attacks. By continuously adjusting thousands of firewall rules per minute, ADAMnetworks ensures **that ZTA remains responsive to emerging risks, regulatory demands, and technological advancements, fulfilling the primary goals of implementing ZTA guidelines.**

ADAMnetworks technology advancement proposes a dynamic policy enforcement approach that reflects the industry's evolution from traditional firewalls, transparent packet inspection, and next-generation Unified Threat Management (UTM) toward more sophisticated threat intelligence feeds that adapt to encrypted traffic. The approach emphasizes **the opportunity to redefine and implement adaptive tools within ZTA guidelines, ensuring that the architecture meets and evolves with the dynamic demands of ZT security, ultimately fostering a more resilient and proactive defense posture.**

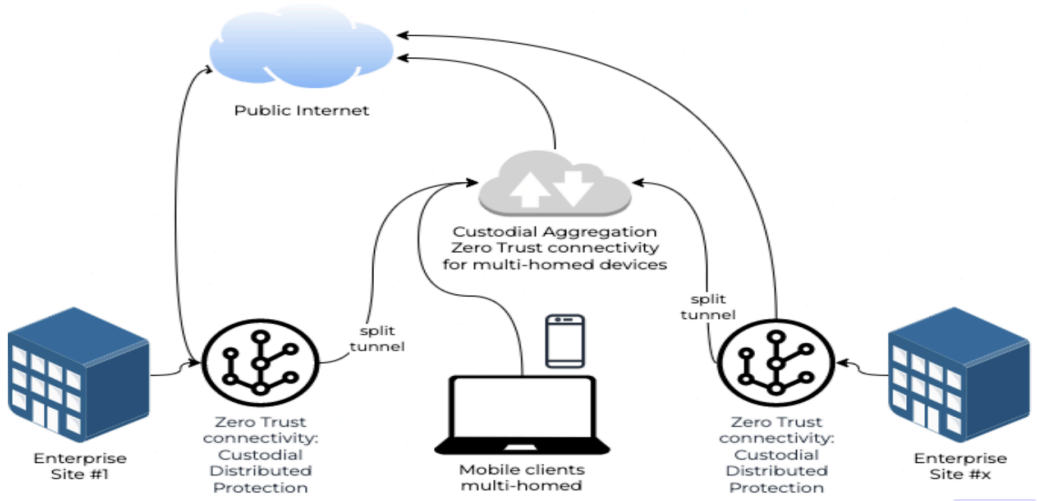

(Honoring TLS and end-to-end encryption without proxies- Source: ADAMnetworks)

4. ZTA Risk-Based Data-Defined Approach

As detailed in subsection 8.4, the implementation of ZTA guidelines should be approached with a risk-based model that prioritizes data value. **Kiteworks stands out in this regard, offering a content-defined ZT model that places a strong emphasis on protecting sensitive data based on its value, regardless of its location.** This unique approach aligns directly with the need to segment infrastructure and implement granular access protection. Kiteworks' ability to enforce **segmentation at multiple levels—application, host, and network—ensures that critical resources are isolated in their own trust zones, protected by Policy Enforcement Points (PEPs), while allowing lower-value resources to share zones**.

By applying risk-based segmentation that limits the potential impact of breaches and makes monitoring traffic easier, Kiteworks offers broad and innovative support for a robust ZTA implementation. Its centralized governance capabilities ensure consistent policy application, enabling organizations to identify and protect each resource with the appropriate Policy Enforcement Point (PEP). By aligning with compliance requirements such as NIST 800-171 and ISO 27001, Kiteworks ensures that stringent access controls and segmentation are maintained, mitigating risks associated with data exfiltration and unauthorized access, which is essential for maintaining an **adaptive and relevant ZTA framework.**

Kiteworks' adaptability is a key feature that directly addresses the need for a phased and evolutionary approach to ZTA implementation, **as described in subsection 2.3's Project Approach.** By offering adaptive policy enforcement and dynamic monitoring capabilities, Kiteworks supports organizations in **incrementally enhancing their security posture without disruption.** It allows organizations to begin with their existing infrastructure and integrate functionalities to gradually evolve toward a more comprehensive ZT model.

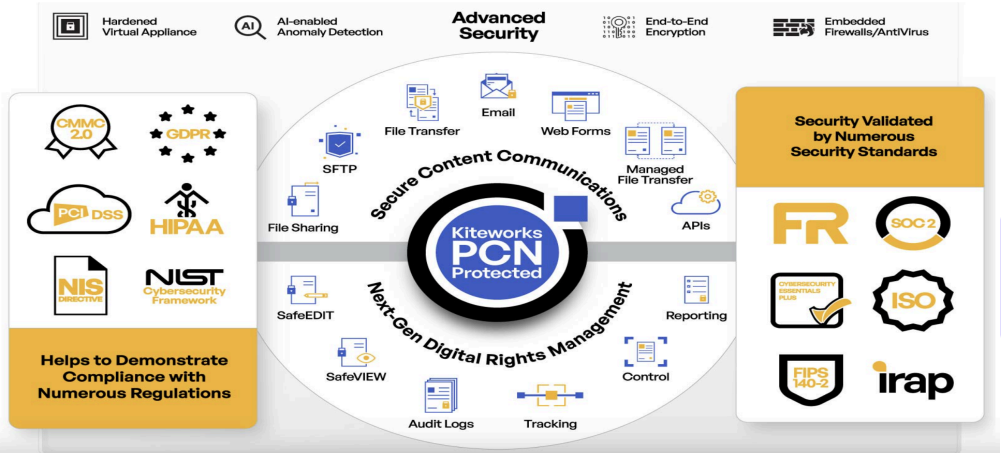

(Kiteworks Unifies, Tracks, Controls, and Secures Sensitive Content Communications- Source: Kiteworks)

As organizations move beyond perimeter-based protections, formulating access policies becomes essential, as subsection 8.2 significantly outlines. Kiteworks uniquely excels in centralizing governance and enforcing consistent access policies. It enables organizations to **adapt access controls in real-time based on factors such as user roles, device types, employment arrangements, and work locations.** This dynamic and responsive nature ensures access policies effectively support secure data exchange across communication channels. Further, Kiteworks's capability aligns with the need for detailed access policies that specify who can access each resource and under what conditions based on user roles, device types, and access requirements. By supporting constraints such as location, time of day, and employment arrangements, Kiteworks ensures that access policies are tailored to interconnected environments, allowing organizations to maintain control while adapting to real-time risk assessments.

Addressing subsection 8.3, which emphasizes integrating and identifying existing security capabilities within a ZT framework, Kiteworks offers unified visibility and comprehensive audit trails, enhancing an organization's compliance posture. Its real-time threat detection capabilities allow organizations to seamlessly leverage and **integrate existing technologies, ensuring that their ZT strategies are effective, adaptable, and aligned with regulatory standards such as NIST 800-171 and ISO 27001.** This capability not only strengthens threat detection and response but also ensures that all data exchanges remain secure and monitored in real-time.

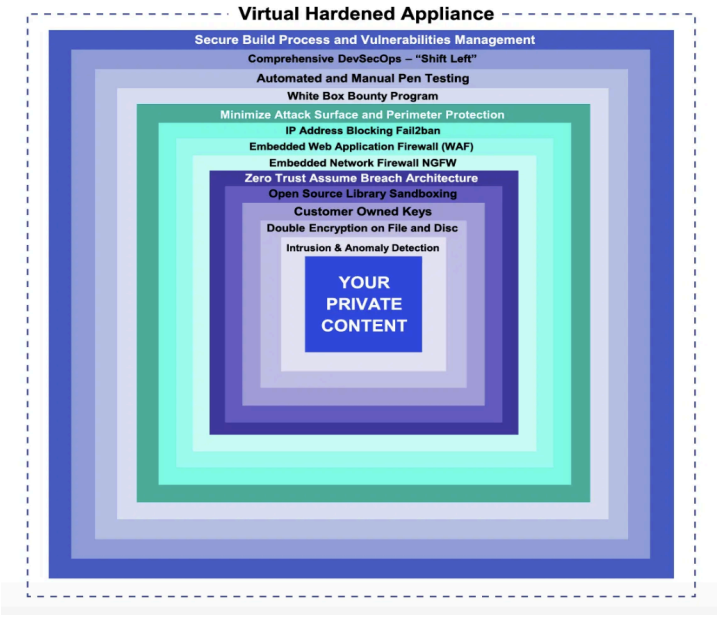

(Kiteworks Secure Managed File Transfer goes beyond protecting sensitive content with encryption protocols and storage- Source: Kiteworks)

Kiteworks further reinforces the principles outlined in subsections 3.1 and 3.4, which stress the importance of secure collaboration and data security within a ZTA framework. Kiteworks facilitates **secure data sharing through encrypted communication channels,** ensuring that sensitive information remains protected and inaccessible to unauthorized entities. Operating within frameworks such as Software-Defined Perimeter (SDP), microsegmentation, and Secure Access Service Edge (SASE), Kiteworks effectively implements a content-defined ZT model, aligning with ZTA principles and preventing unauthorized access to critical data assets.

With unified visibility, comprehensive audit trails, and real-time threat detection, Kiteworks provides a cohesive, adaptive, and efficient solution that supports regulatory compliance and meets the evolving needs of a ZTA framework. **It ensures that organizations can address the complexities and challenges of the evolving cybersecurity landscape while meeting the ZTA's required data protection and compliance standards.**

As the NIST SP 1800-35 implementation evolves, we look forward to discussing these proposals with NIST and are available for any questions. We remain excited to collaborate with NIST to increase engagement with innovative companies.

Respectively,

/s/ Michelle Sahar
Michelle Sahar
Cybersecurity Policy Director, OpenPolicy