Zero-Trust But Verify Critical Infrastructure Systems Using a Known Adversary Security Engine

Kyle Sullivan, *Member, IEEE*

***Abstract*—As industrial systems are connected into cloud environments to leverage the benefits of powerful modern computing and artificial intelligence technologies, we will witness a drastic change to the cyber-physical security frameworks that have traditionally relied on physical isolation to mitigate the security risks caused by the implicit trust needed to ensure real-time guarantees. The security implications introduced in the new era of cloud-powered industrial systems are discussed in this paper and a novel zero-trust based security model is introduced to address the security challenges. The model is based on emerging zero-trust architectures combined with the principles of a "Known Adversary Security Engine" to "zero-trust but verify" the security of critical infrastructure systems from the perspective of a trusted known adversary embedded within the network. The model presents solutions to the security challenges discussed by leveraging the zero-trust principles of micro-segmentation, least privileged access, and "never trust, always verify" integrated into continuous security monitoring. Together the zero-trust powered security integrated with real-time perspectives from an embedded known adversary forms a robust security framework to address some of the concerns introduced by cloud-computing into the physical realm of industrial systems. The effectiveness of the proposed model is further examined through a case study to illustrate how the new framework enhances the security posture in the evolving field of cyber-physical security (CPS). Finally, the paper provides guidance on future research to further advance the rapidly evolving field of CPS.**

***Index Terms*—Zero-Trust Architecture, Critical Infrastructure Systems (CIS), Known Adversary Security Engine, Cyber-Physical Security (CPS), Cloud Computing Security, Fog Computing, Network Security, Micro-Segmentation, Continuous Authentication, Dynamic Policy Enforcement, Threat Detection, Resilience, Industrial Control Systems (ICS), Real-Time Security, Security Frameworks**

I. Introduction

THERE exists a world of relatively unknown hardware and embedded devices that control the critical infrastructure powering modern society. Found within Industrial Cyber-Physical Systems (ICPS), these industrial devices blend together cyber and physical domains to form fulfill a critical technological role. Without these types of devices, the systems that uphold important industries such as manufacturing, communications, utilities, energy, and public safety would not exist as they do today. These industrial systems depend on secure real-time operations to provide essential services without interruption. But with the introduction of newly emerging cloud technologies, these little known yet vital pieces of technology face new security threats that could have widespread damaging impacts to our society.

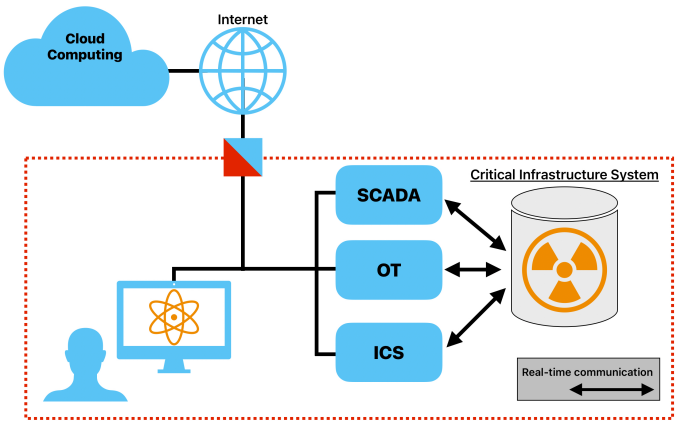

Fig. 1. Overview of a Cloud enabled CIS network environment.

A. Background

Historically industrial systems have lagged behind the cutting-edge of computer security technology due to organizations often treating them as an edge case or an exception to policy. [2] Largely the reasons for the special treatment of these systems were based on practical purposes, often due to their physical nature of these systems and the real-time role they perform. Industrial systems typically have critical real-time requirements which demand high reliability and enforce strict safety guarantees; simply put failure is not an option.

But with the rapid introduction (and convenience) of new cloud based technologies, the traditionally ignored "exceptions to security policy" found within industrial systems are being thrust into the modern era of cloud computing; creating a future era of security vulnerabilities.

B. Relevant Concepts

Tackling emerging security challenges can sometimes seem difficult to conceptualize. Before diving into the problems impacting cloud based CIS and examining a solution its helpful to review relevant concepts field of cyber-physical systems.

1) Critical Industrial Systems: While all critical industrial systems involve cyber-physical connections making them a type of ICPS, CIS are specific types of cyber-physical systems that power and control vital aspects of society. CIS are often found supporting critical sectors responsible for the functioning of an economy such as communications, energy, transportation, or water. These important sectors are usually overseen by government agencies such as the Federal

Communications Commission (FCC), Department of Energy (DOE), Department of Transportation (DOT), the Department of Homeland Security (DHS), or equivalent national level agencies. Typically the agencies that govern these critical industries impose strict regulations and rules to ensure the safety and reliability of CIS which introduces unique challenges as compared to standard commercial systems. These industrial systems are often implemented through Supervisory Control and Data Acquisition (SCADA) devices, Operational Technology (OT), and Industrial Control Systems (ICS). Collectively these technologies and their related architectures will be explored in this paper as CIS.

2) Cloud and Foggy Computing: In the context of Industrial Cyber-Physical Systems (ICPS) and Critical Infrastructure Systems (CIS), cloud computing introduces new capabilities to super-charge the processing power by leveraging the vast computing and storage resources of the cloud. Cloud computing in CIS also allows for scalability and flexibility of resources across a traditionally physically limited environment. Cloud computing also improves operational efficiency by distributing computational workloads across cloud nodes. The concept of cloud computing in CIS is further extended through the use of "foggy cloud computing", which introduces Edge computing closer or even within the physical networks of industrial systems. Foggy computing improves on the classic cloud computing design by localizing processing power physically near the CIS embedded devices to allow for more real-time responsiveness and reduced latency. Foggy computing is an important concept needed to support the real-time demands of CIS and their safety guarantees. Foggy computing also provides improved security over cloud computing by processing data more localized to the secure industrial networks and reducing the need to transmit potentially sensitive data across public cloud infrastructure. Overall, foggy computing is an ideal solution for cloud-computing in ICPS and CIS environments due to its ability to meet real-time requirements and balancing efficiency with safety. Throughout this paper cloud computing will be discussed with foggy computing as the ideal use case for CIS networks.

3) Zero-Trust Architecture: Zero-trust is an emerging security framework designed to address a major cause of cyber breaches, lateral movement. Traditional cybersecurity frameworks approach security by hardening networks against external threats through the use of firewalls and access control tools to prevent ingress/egress into a network. This approach has fallen short over recent years as malicious actors leverage the implicit trust that traditional networks grant to internal devices and connections originating from inside of a network. Malicious actors have leveraged stolen or breached credentials to escalate their privilege from within a network and bypass the hardened outer defenses. Once inside a breached network, the malicious actors move laterally across the network and can pivot into new devices or user accounts allowing them to penetrate deeper into a vulnerable network. Zero-trust architectures defeat lateral movement and privilege escalation by denying access to resources and devices explicitly and enforcing continuous authentication throughout the network. In doing so, network devices must authenticate their connections with a zero trust policy engine which enforces the zero trust rules across the network. The policy engine has the responsibility of verifying the interactions of all devices and users on the network and in doing so uphold the principles of zero-trust.

4) Trusted Known Adversary: One of the founding ideas behind zero-trust architecture is the idea of "never trust, always verify" which leads to the concept of "assume breach". This means that IT organizations implementing zero-trust architectures should always assume there is a malicious actor (or adversary) on their network and utilize a zero-trust engine that will isolate malicious activity and prevent any lateral movement. This mentality fully embodies the zero-trust principles and is the core idea behind the use of a "trusted known adversary". In order to fully verify zero-trust, security engines should incorporate the insights and perspectives of a trusted known adversary on the network to provide real-time feedback about what a malicious actor would be able to see if there was an active security breach on the network. These insights provide invaluable data for zero-trust security tools (SIEM/SOAR) and supports the continuous monitoring and automated response that is required to fully secure modern cloud-based networks. A trusted known adversary implemented on a zero-trust network and integrated with security tools also facilitates automated adversary emulation which will enable a robust and real-time security monitoring and response ecosystem.

C. Objective and Goals

Throughout this research we will be examining security challenges that face industrial systems in the cloud era, specifically a class of Industrial Cyber-Physical Systems (ICPS) known as Critical Infrastructure Systems (CIS). While the proposed security model and following discussion presented below applies to ICPS as a whole, this paper focuses on Critical Infrastructure Systems as the target use case for the proposed solutions. The goal of this paper is to introduce a novel security model based on the concepts of a zero-trust architecture powered by a Known Adversary Security Engine and its application to cloud based CIS security.

II. Security Challenges of Cloud based CIS

The move from legacy environments which are traditionally isolated into cloud enabled platforms introduces new classes of security challenges and vulnerabilities. The challenges arise when connecting the physical real-time dependent industrial devices into a complex cloud connected environment. Traditional CIS devices often rely on the inherent trust of their physical nature with an emphasis on real-time guarantees which leaves the door open for a new class of cloud based vulnerabilities that did not exist when they were originally designed.

A. Increased Attack Surface

By departing from the air-gapped model that traditional industrial systems often relied on, the attack surface of

CIS environments has grown beyond physical isolation of network devices. With the integration of cloud computing into industrial networks, CIS devices now exist in a more complex and interconnected system exposing them to more than the physical threats but to a new emerging class of cloud based attacks. With access to cloud services and potentially third-party external APIs, malicious actors can more easily find access points into the networks supporting our critical infrastructure.

B. Lateral Movement

A typical goal of an adversary once they gain access into a compromised device is to attempt to pivot across the network laterally into additional devices or systems. This type of attack is most successful in networks that have inherent trust between internal devices which is common in traditional industrial systems. The cyber-physical nature of industrial networks typically results in little to no authentication between devices which makes them vulnerable to lateral movement. Lateral movement is a serious threat posed by adversaries that if successful can result in significant damage, but can be thwarted by network segmentation and continuous monitoring of unauthorized activity. Additionally, zero-trust architectures also implement policies that enforce strict authentication between devices and can prevent lateral movement type of malicious activity from spreading across the network.

C. Misconfigurations Due to Increasingly Complex Environments

The complexity of emerging cloud environments increases the potential for configuration errors. These misconfigurations can be lead to vulnerabilities that could be leveraged by a malicious actor. Cloud enabled industrial networks could be implemented through multi-cloud or hybrid platforms which introduces additional risks of misconfiguration or inadequte security controls due to the complex environment. There is also a supply-chain risk that is introduced, either through reliance on cloud updates or even third-party vendor dependencies that can introduce their own class of vulnerabilties.

D. Zero Trust Enforcement Gaps

Properly implemented zero-trust architectures are designed to minimize the implicit trust that exists in legacy networks, often by requiring continuous authentication of users and devices followed by authorizing interactions and connections. However, with zero-trust technologies being relatively new it could prove difficult to implement the policies properly in an industrial network due to the legacy devices and unsupported protocols. This gap in security could lead to scenerios where zero-trust policies are not properly enforced which could potentially leave important parts of the network vulnerable. These types of misalignment between the policies set by the zero-trust engine and the faulty implementation in CIS devices could introduce new security concerns and potential vulnerabilties.

III. Introducing the Known Adversary Security Engine: A Security Use KASE

The challenges brought about by cloud-based CIS environments require a novel approach to address the new vulnerabilities threatening cyber-physical security. Combining the benefits of zero-trust architectures such as micro-segmentation, authentication and authorization, and continuous monitoring with adversarial emulation and real-time threat detection/mitigation we can develop a new comprehensive security model. This new model implemented alongside of the zero-trust policy engine is referred to as the "known adversary security engine". This framework addresses the gaps that exist when implementing zero-trust technologies inside of a CIS environment, ensuring that zero-trust is verified everywhere.

A. Key Features of the Security Engine

An important feature of the KASE security model is a tight integration between real-time feedback and security tools, to include Security Information and Event Management (SIEM) and Security Orchestration, Automation, and Response (SOAR) tools.

1) SIEM/SOAR API Integration: In order to fully support security operations within a zero-trust environment it is important that the security engine is able to interface with a wide variety of security tools. Currently, SIEM/SOAR tools implement proprietary or disjoint APIs which would be overcome by an orchestration layer inside of the security engine. This orchestration would be responsible for ensuring that security findings can be appropriately reported and cordinated with the complex suite of security tools often found within modern networks, especially in cloud based environments.

2) Adversarial Emulation: The security engine would be responsible for using the latest tactics, techniques, and procedures (TTPs) to emulate current and emerging threats from potential adversaries. By leveraging adversarial TTPs against network devices from within the network, the security engine would gain valuable insights of the threats that actively threaten the network in real-time. The security engine would able be integrated with the zero-trust policy engine to verify zero-trust policies are being followed properly. Importantly, any security or policy violations can be reported and coordinated with the suite of security (SIEM/SOAR) tools to assist with a dynamic security response.

3) Verifying Zero-Trust: Perhaps the least thought of but most impactful purpose of the "known adversary security engine" model would be to verify zero-trust. In a true zero-trust architecture, you must assume that there is a security breach. The security engine would play the role of "assumed breach" and provide feedback to the zero-trust architecture that devices are following zero-trust policies appropriately. The security engine can be thought of as a "good guy on the inside" that is playing the role of an adversary but feeding back results to the proper authorities.

IV. KASE Study: Adversarial Insights of Mission Critical Systems

The Known Adversary Security Engine is a powerful zero-trust security framework that can enhance the security of

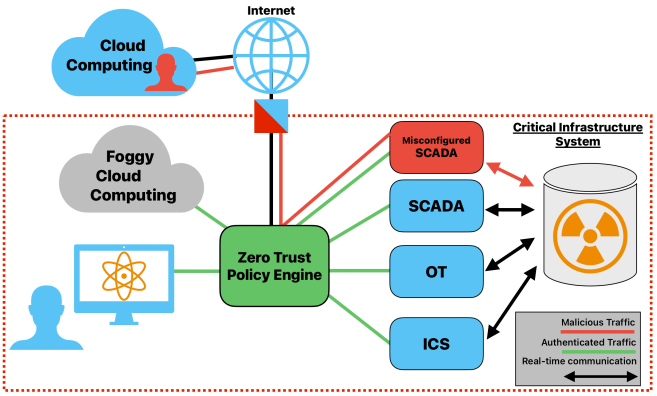

Fig. 2. Overview of a Scenerio with Malicious Actor Abusing Misconfigured Devices.

CIS embedded devices and provide real-time insights into the security impacts of a potential adversary that is assumed to be present and active in the network.

A. Scenerio Overview

The hypothetical scenario we will examine consists of critical infrastructure systems controlling a Nuclear Power Plant which are connected to a foggy cloud computing environment secured by zero-trust architecture.

In this scenario, updated environmental policies from the government's environmental agency requires the Power Plant to scale their cloud-computing resources so they can provide additional data to inspectors. At the same time, new department of energy regulations emerge which requires additional SCADA devices to be installed at the plant to address safety concerns and provide redundancy for the real-time systems. The competing priorities to meet their regulatory requirements from government agencies causes the network engineers and IT personnel at the power plant to work disjointly which results in misconfigured SCADA devices being added to the network and zero-trust policies not being enforced properly.

B. Problems Arise at the Plant

Unknown to the personnel, the new SCADA devices are now communicating with the publicly routable cloud-based environment agency API. This unintentional vulnerability allows outside traffic into the network on commonly used network ports as long as it originates from the same public cloud that hosts the environmental API endpoint. Because the zero-trust policy engine was misconfigured and is allowing any traffic originating from public cloud hosting the environmental API, malicious actors discover they can reach the new SCADA devices installed internally to the power plant network. Armed with this unauthorized access, the malicious actors launch an attack on the power plant and damage critical systems forcing the power plant operators to shut down the entire plant to prevent a meltdown from taking place. The sudden loss of power cascades rapidly across the regional power grid severely disrupting large population centers and results in a national emergency being declared. Disaster has struck.

C. Learning from The Enemy Before They Attack

Now let's examine the same scenario only this time the power plant network implemented a Known Adversary Security Engine to actively verify zero-trust.

Operating continuous within the network are trusted known adversary devices that are synchronized with the zero-trust policy engine and emulate adversarial techniques to verify zero-trust policies with real-time feedback integrated into the security tools (SIEM/SOAR). Collectively this forms the Known Adversary Security Engine which is actively verifying zero-trust policies on the network from the perspective of a malicious actor. This enables automated adversary emulation with feedback provided to the security tools for real-time response. Not only are the CIS devices on the network following zero-trust policies but those same policies are being verified continuously from the perspective of an adversary.

In the new scenario, when the misconfigured SCADA devices are introduced onto the network the KASE devices actively detect and start to verify that the devices are following zero-trust policies. The KASE device identifies that one of the new SCADA devices is accepting inbound network traffic on a default port without first receiving permission from the zero-trust policy engine. Normally, the zero-trust PE would verify that the inbound connection is from the proper cloud endpoint hosting the legitimate environmental API but because both the device and PE are misconfigured it allows inbound traffic from any endpoint in the public cloud. Luckily the KASE device is able to connect to the SCADA device without authenticating its traffic. KASE reports in real-time that the SCADA device is not following zero-trust policies and informs the security tools. The policy violation information is fed automatically to the network security tools integrated into KASE. The security tools can then automatically quarantine traffic by isolating the micro-segmented network that the misconfigured SCADA device belongs to while also notifying security operations personnel. Not only does this automated response prevent the security breach from happening in the first place but it also provides valuable insights to the security team to address the situation. They are able to review the security logs reported from the KASE device and realize they need to reconfigure the SCADA device to first authenticate with the zero-trust policy engine when interacting with the newly implemented cloud-based environmental API. With the solution in place, the properly configured SCADA devices now authenticates with the policy engine which verifies the inbound connections are coming from the legitimate environmental API directly. The environmental inspectors are happy with their new data, the energy regulators are satisfied with the improved redundancy, and a meltdown has been avoided.

D. Lessons Learned

This case study highlights the importance of verifying zero-trust principles from the perspective of a known adversary. Implementing zero-trust is only one piece of the overall solution. Having a trusted actor embedded into zero-trust architecture which performs the role of an adversary enables the network to "always assume breach". To fully implement "never trust,

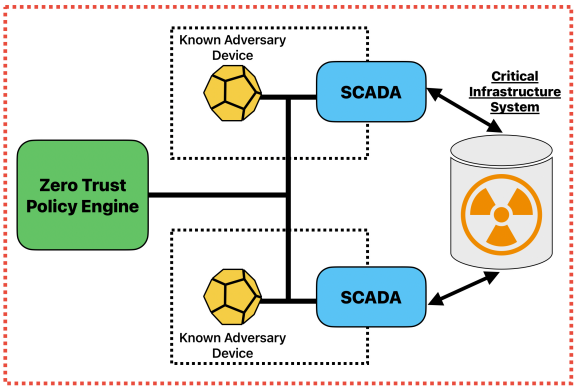

Fig. 3. Zero-Trust Implementation with Micro-Segmentation and Known Adversary devices.

always verify" then a robust security framework should also be "never trust, always verify" the zero-trust policy engine itself through the perspective of a trusted known adversary.

V. Future Research Guidance

Zero-trust is an emerging technology which will continue to evolve as it is implemented across more and more networks. At the same time, with the growing introduction of ICPS and CIS devices into cloud-environments it is a natural progression to implement zero-trust in the most critical of physical networks. The principles of a trusted known adversary in this paper are novel in nature and warrants additional research into the technical implementation.

A. Security Tool Convergence

One area of importance is the convergence security tools and a unified API or data format for reporting of security events. A current barrier to fully implementing a known adversary security engine is the wide variety of security tools that currently exist in the industry with differing APIs, protocols, and data formats.

B. Universal Adversary TTPs

Another area to explore is a common format for sharing adversary TTPs and techniques. In order to have a fully effective security engine with realistic adversary emulation there will need to exist relevant and accurate data on the latest adversarial TTPs, techniques, and indicators of compromise. A properly implemented KASE should not only verify zero-trust policies on a network but it should do so while emulating adversarial techniques. In this way, realistic insights of a malicious actor can be provided to automated security tools.

C. Embedded Trusted Adversary Devices

Industrial systems have unique implementation challenges due to their physical nature and real-time requirements. An additional area to research would be into compatible embedded devices which can implement the principles of adversary emulation while also matching the form factor and architecture of ICPS and CIS devices.

VI. Conclusion: Making the KASE for Security

In this paper, we introduced a novel security model to address emerging threats impacting cloud enabled critical infrastructure systems. Through our case study, we examined how security challenges brought on by cloud-enabled CIS environments can be leveraged by an attacker to achieve catastrophic results. We then demonstrated how using a novel security model based on zero-trust powered by adversarial insights can prevent an attack before it happens.

While zero-trust architectures are a strong step forward to securing cloud environments, there are still security gaps that exist potentially threatening the devices that power our critical infrastructure. The use of a "known adversary security engine" attempts to close these gaps by providing orchestration of adversary emulation with real-time response from security tools. Furthermore, the no-fail requirements of CIS devices warrants the need to verify zero-trust which is accomplished through the insights of a trusted known adversary existing within the network. Because something as simple as a policy misconfigurations can open the door to malicious actors, it is vital that our critical infrastructure is protected by a robust and comprehensive security model capable of addressing problems in real-time. Because the outcome of failure is potentially catastrophic, security is more important than ever. With concepts like the "known adversary security engine" we can ensure that cloud enabled industrial cyber-physical systems are continuously protected from the threats posed by rapidly emerging cloud technologies.

Acknowledgments

The author would like to thank Dr. Andy Novocin, Associate Professor of Computer Science at the University of Delaware. His mentorship has aided in my career journey, and his dedication to teaching and advancing the field of Computer Science served as an inspiration for this research.

References

[1] "Advancing Zero Trust Maturity Throughout the Network and Environment Pillar," National Security Agency, https://media.defense.gov/2024/Mar/05/2003405462/-1/-1/0/CSI-ZERO-TRUST-NETWORK-ENVIRONMENT-PILLAR.PDF

[2] S. Rose et al., "Zero Trust Architecture - NIST Technical Series Publications," National Institute of Standards & Technology, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf

[3] "DoD Zero Trust Strategy," Department of Defense, https://dodcio.defense.gov/Portals/0/Documents/Library/DoD-ZTStrategy.pdf

[4] "Zero Trust Maturity Model Version 2.0," Cybersecurity and Infrastructure Security Agency, https://www.cisa.gov/sites/default/files/2023-04/zero_trust_maturity_model_v2_508.pdf

[5] "Department of Defense Zero Trust Reference Architecture," Department of Defense, https://dodcio.defense.gov/Portals/0/Documents/Library/(U)ZT_RA_v2.0(U)_Sep22.pdf

[6] J. R. Biden, "Executive Order 14028 on Improving the Nation's Cybersecurity," The White House, https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/

VII. Biography

Kyle Sullivan (member, IEEE) received the B.S. in computer science (cum laude) (2016) and the M.S. in cybersecurity (2021) both from the University of Delaware. Currently he is a Field Engineer working in the defense industry in Columbia, Maryland. He also is a Major in the U.S. Army National Guard where he serves as a Network Engineer supporting military communication and network systems. His research interests include the security of real-time systems and embedded operating systems. He has also worked on various defense projects related to the security and hardening of forward deployed systems with a focus on improving their security and reliability. He previously has published an article in a cybersecurity journal related to securing network and communication systems against adversaries. Mr. Sullivan is also a member of the Military Cyber Professionals Association.