



ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

A Zero Trust approach for the cybersecurity of Industrial Control Systems

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

A Zero Trust approach for the cybersecurity of Industrial Control Systems / Zanasi, Claudio; Magnanini, Federico; Russo, Silvio; Colajanni, Michele. - ELETTRONICO. - 21:(2022), pp. 1-7. (Intervento presentato al convegno The 21st IEEE International Symposium on Network Computing and Applications (NCA 2022) tenutosi a Boston nel 14-16 December 2022) [10.1109/NCA57778.2022.10013559].

This version is available at: <https://hdl.handle.net/11585/954287> since: 2024-01-29

Published:

DOI: <http://doi.org/10.1109/NCA57778.2022.10013559>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

(Article begins on next page)

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

A Zero Trust approach for the cybersecurity of Industrial Control Systems

Claudio Zanasi*, Federico Magnanini[†], Silvio Russo*, Michele Colajanni*

*Department of Computer Science and Engineering, University of Bologna, Italy

{claudio.zanasi4, silvio.russo2, michele.colajanni}@unibo.it

[†]Department of Engineering “Enzo Ferrari”, University of Modena and Reggio Emilia

federico.magnanini@unimore.it

Abstract—Industrial plants are adopting an increasing number of digital interconnected technologies that are enriched by several software applications. The IT/OT convergence offers several benefits in terms of efficiency and flexibility but it opens as many issues in terms of cyber vulnerabilities because industrial plants were not designed to be open to Internet. The frequency of successful cyber attacks shows that typical security solutions are inadequate to the novel complexity of industrial contexts. This novel scenario requires original approaches differing from traditional multi-layer networking solutions that are applicable just to rigid and stable infrastructures. We explore the applicability of Zero Trust Architecture (ZTA) principles to the industrial context by designing, implementing and testing an integrated defensive solution. The results obtained through a working prototype show that it is possible to implement a Zero Trust identity-centric approach in an industrial context to increase the security and flexibility of the system while providing complete visibility over the entire network. The proposed approach can be used to strengthen legacy industrial systems that were designed for offline use, and to allow the adoption of innovative technologies that minimize the cyber risk for the overall infrastructure.

Index Terms—cybersecurity, networks, zero trust, NGFW

I. INTRODUCTION

The rapid adoption of advanced process control systems such as the Industrial Internet of Things (IIoT) and Industry 4.0 is completely changing the security requirements for cyberphysical and industrial systems. The need for remote accesses, including smart working and remote maintenance, the evolution of supply chain workflows, and the acceptance of the Bring Your Own Device (BYOD) model are exposing the limitations of perimeter-based defense strategies. In addition, the impressive and increasing number of successful cyber attacks confirms the need for novel security approaches.

The previous obsolete defenses were oriented to assume that, after the design of a security perimeter, everything inside it was trusted implicitly. In fact, the complexity of modern systems and the need for network connections, even for industrial plants, require a complete re-design of *trust* given to each specific component and user of the system. Zero Trust is a model where different elements of the architecture do not implicitly trust each other, but there is a continuous identification and authorization process for all the components. The Zero Trust security model is a set of principles and strategies based on the awareness that cyber threats exist

and can originate from outside but also inside the perimeter of traditional networks. Google originally implemented this approach by releasing their internal security framework BeyondCorp [25]. Zero Trust is becoming an emerging paradigm shift for cybersecurity that is being adopted by the most mature actors. The main principles have been formalized in a NIST document [22] that describes the core aspects of a Zero Trust Architecture (ZTA). Also, the US government has recently published a memorandum [7] that makes the adoption of this model mandatory for all internal agencies. The Zero Trust security model is a set of principles, not a specific architecture. Hence, everything must be tailored to the specific needs and goals of the infrastructure to which we want to apply it. Even more importantly, it requires a change of mindset that security managers and architects must adopt.

We think that when dealing with cyber physical critical infrastructure and industrial control systems (ICS), the Zero Trust model should become an essential method that can prevent, detect and react to the increasing number of cyber attacks. The *castle-and-moat* model is dead, and perimeter-based security solutions can no longer protect infrastructures from threats from inside and outside the network perimeter [14]. In addition, the security model of industrial systems relied heavily on the idea of security-by-obscurity by using proprietary closed-source protocols and physical isolation from IT networks to limit the attack surface area.

Today, with the convergence of IT and ICS architecture [17] [13], this security model is no longer valid, thus exposing the weaknesses of these systems and making them among the primary victims of cyber-attacks. It became clear that for ICS systems, the real question is not “if” an intrusion will take place but “when”.

The solution proposed in this work results from a detailed analysis of the state-of-the-art best practices for the security of the industrial sector. The established Industrial Demilitarized Zone (I-DMZ) [12] model allowed to defend this type of infrastructures. With multiple interconnections, edge- and cloud-based architectures, remote maintenance operations I-DMZ solutions are becoming obsolete because they do not guarantee protection and flexibility of modern industrial plants. This paper introduces an innovative Zero Trust Architecture for the industrial sector based on a robust continuous Authentication and Authorization system and a per-device segmentation. We

demonstrate that a modern cybersecurity architecture can improve the defense capabilities of IT-oriented industrial systems. The architecture is based on an identity-centred security approach [21] and guarantees complete visibility over the entire network. All the authorizations are based on a continuous evaluation of the risk level and the context in which the connection is made. Furthermore, it is based on a trust algorithm that considers not only the user information but also the device identity and its health. The trust algorithm is the core component that gives the necessary flexibility to our solution. It is dynamic and capable of adapting to the architecture's needs and security posture by considering several different parameters and the information the company deems most appropriate to authenticate and authorize its employees.

The proposed solution has the advantage of being suitable to both on-premise and cloud-based architectures, where Big Data and AI applications of industrial data are externalized to third-party cloud providers. On the other hand, similar flexibility would require substantial modifications of a multi-layer network architecture, such as an I-DMZ, to the extent of making it impractical.

The remaining part of this paper is organized as follows. Section II discusses the related work and highlights our contribution to the state of the art. Section III describes our approach and the proposed architecture. Section IV describes our prototype implementation used for experimental evaluation and a brief comparison with the I-DMZ approach. Finally, Section V concludes the paper.

II. RELATED WORKS

The core principles of the Zero Trust model are described in the review by Syed, Naeem Firdous, et al. [23]. This paper presents the best practices for the adoption of each *security tenet* as identified by NIST [22] along with a comparison of possible implementation strategies. As important contributions, it highlights main weaknesses of current approaches, points some areas of improvement and focuses on the ideal context for the implementation of each security practice.

While the Zero Trust model is rapidly gaining adoption in the IT sector, in the industrial context it is much harder to introduce innovative security practices. The current approach to industrial cybersecurity is to use a layered defense strategy as promoted by international security standards, such as ISA/IEC 62443 [10]. Due to the ever increasing needs for network connectivity of industrial systems the concept of industrial DMZ (I-DMZ) has been introduced after several decades of the use of similar architectures in IT systems. The I-DMZ is a perimeter network placed between Informational Technology (IT) and Operational Technology (OT) zones that acts as an intermediary for any communication between the two network segments. This solution can be considered a partial implementation of the Defense-in-Depth Framework, which has as its primary goal the protection through isolation of the OT network. There are numerous examples concerning I-DMZ based solutions. For example, Jiang et al. [11] propose an architecture with paired firewalls to better enforce the

separation of the IT side and the OT side by placing a dedicated firewall to each leg of the network. This scenario allows the implementation of stricter security rules with respect to a configuration with a single firewall while simplifying the implementation and the management of the infrastructure thanks to a separation of the responsibilities. The same authors investigate the performance implication of this paired firewall I-DMZ architecture [12]. Alvarez, et al. [3] propose the design and implementation of an I-DMZ for a real use case to protect highly critical real-time data produced by a manufacturing plant during its operations. A test was successfully performed to validate the security proprieties of the architecture and the conformity to the initial functional requirements.

In industrial contexts, the Zero Trust security model is still at an early adoption phase, hence few contributions exist in literature. For example, Algappan et al. [2] present an assessment on cybersecurity posture of distributed energy resources and virtual power plants. These components are facing unique cyber threats due to their strategical importance and to the direct interconnection to the power grid. Hence it is critical to ensure a proper cybersecurity posture. Instead of a defense-in-depth approach, this work evaluates the benefits of a defense strategy based on Zero Trust Architectures. The intrinsic decentralization of these systems makes it difficult to define appropriate organizational perimeters hence a Zero-Trust approach can provide better resilience to the overall system in case of successful attacks. Another related work is proposed by Mir and Ram Kumar [16], where the authors analyze the application of ZTA in a conventional industrial setting. They include an overview of the security improvement offered by the Zero Trust model with respect to existing best practices, some key elements for a successful implementation and a high level view of a possible implementation.

These two papers [2], [16] are pioneering to consider Zero Trust security models in an industrial context, but they offer just a high level reference architecture together with a theoretical discussion about possible benefits without providing experimental results. We propose a novel Zero Trust Architecture that is tailored for an industrial setting based on access proxy and next generation firewall technologies. We implement a proof of concept and test successfully both solutions. We also validate the security capabilities of the solutions by simulating a realistic scenario composed by an industrial network with a compromised host trying to perform lateral movements and disrupt the other connected systems. We demonstrate that the adopted security measures are able to successfully stop multiple attack attempts thus showing the defensive capabilities of the implemented ZTA.

III. PROPOSED ARCHITECTURE

This section describes our proposal of a Zero Trust Architecture for the industrial sector. It overcomes many limitations of the traditional I-DMZ solution, allowing at the same time the implementation of new services and functions that the industrial sector needs without compromising its security operations and safety. It is based on a novel approach with




Fig. 1. Proposed architecture

respect to the industrial system context. It avoids any concept of trusted internal perimeter thanks to a strong authentication and authorization mechanism, allowing a sysadmin to monitor and control any communication inside the network.

By considering the possible implementation strategies for ZTA described by the NIST [22], our solution can be classified as an hybrid between *ZTA Using Enhanced Identity Governance* and *ZTA Using Micro-segmentation*. We use strong access management and user/device identity governance. We also introduce a strong segmentation in the network to protect the critical devices of the OT from any potential threats. Although the classical Micro-segmentation implementation is largely adopted, it requires the concept of a perimeter. Even if it is small, it implies the concept of trust inside it thus being in contrast with the basic concepts of the Zero Trust Approach. For this reason, we propose a per-device network segmentation. Only through this solution, we can monitor and manage any interaction between any device.

The Architecture is shown in Figure 1. It is based on two fundamental components: the Access Proxy (AP) and the Next Generation Firewall (NGFW). They enable the implementation of fine-grained Authentication and Authorization processes and segmentation of the network to guarantee that the communications, both from outside and inside, can be monitored and controlled. The network can be divided into three main areas: Access Control, IT, and OT. The different areas are connected by the NGFW, which plays the role of Policy Enforcement Point (PEP) and Policy Administrator (PA) to monitor and control any communication in the network. Thanks to the Zero Trust approach and the fine-grained segmentation, our solution is designed to minimize the possibility of lateral movement and external intrusion.

A. Access Control

The first part manages the access to the network; the core component is the Access Proxy which provides authorization and authentication per access. It collaborates with several services to store all the needed data. The Domain Controller (DC) stores all the information about users, their role in the company, and the associated privilege level. The Device Manager (DM) is the equivalent of the Domain Controller for

the devices, and it stores all the information needed to identify the devices, their certificates, and their owner. The SIEM monitors all network events, correlates the logs, and guarantees real-time analysis that, thanks to threat intelligence services, can emit alerts and alarms. This part of the architecture can be extended based on the security needs identified by the organization. It can be enriched by security services that can make the authentication and authorization process more flexible and/or robust. The combination of all these services grant highly secure access to the network.

1) *Single Sign On*: User authentication is a core part of our implementation and security strategy, hence we suggest the use of a Single Sign-On (SSO) system to simplify the access to any company resource. It is up to the policy engine to check the user's role and authorize or deny each request to a resource according to the security policies defined by the organization.

A centralized authentication system helps to create a robust access system, as discussed also in [15]. Often companies use several different authentication mechanisms requiring multiple credentials. This approach does not increase the overall level of security. Instead it could become a security vulnerability leading to weaker credentials. With a single sign-on solution we improve both the security and the user experience of the system. The SSO system is directly connected with the Domain Controller that stores all the information about the user, such as credentials, the role inside the company, and all additional information needed by the policy engine.

The authentication is at least a 2FA to ensure the proper level of security. When the user has provided his 2FA credential, the SSO server contacts the Domain Controller to validate these credentials. If the user is recognized, then the SSO generates a token with all the valuable information for the resource to determine the user's access level.

2) *Access Proxy*: An essential point for the security of any infrastructure is the management of employee's accesses especially remotely [8]. The typical solution is to use a VPN to create an encrypted tunnel between the remote device and the industrial network, but this method will be replaced by solutions based on an Access Proxy. Indeed, a major problem with traditional VPNs is that they are static approach. This means that they allow to grant access to the organization's network, but cannot manage per-service requests. As highlighted in [26], the pandemic led to massive increase in VPN usage for remote work, As a consequence, it became one of the first target of cyber attacks that aimed to obtain a first foothold within the victim's network (e.g., [6] [5] [24]).

The Access Proxy implements the role of Policy Engine. It provides authorization and authentication for users and devices based on fine-grained dynamic policies that take into account the user, the devices information and also the context of the request such as moment of the day or the general risk level of the action to perform. With this solution there is no direct connection between users and applications. Instead, all accesses are routed through the AP which can perform dynamic access control decisions. For example it can restrict

parts of an application that requires a different privilege level. The core component of the AP is the Policy Engine. It represents a fundamental process of a ZTA that executes the Trust Algorithm and it is responsible for the ultimate decision to grant or deny the access to a resource.

The AP is connected to the SSO system, which has the task of authenticating the user by using the information stored in the Domain Controller, and to the Device Manager responsible for the authentication and authorization of devices. It also plays a crucial role in the entire infrastructure and can be considered a single point of failure. However, in a real scenario the AP can be horizontally replicated to improve its availability.

B. Next Generation Firewall

The Next Generation Firewall (NGFW) represents the other core component of our implementation. This firewall provides many security functions in addition to the standard stateful inspection and access decisions. It can block advanced malware and integrate several higher level features like Intrusion Prevention/Detection, Threat intelligence across users, hosts, networks and machines in one device. It can dynamically adapt to the system's state and react fast if something goes wrong thus reducing the overall incident response time.

A NGFW is a powerful device that implements a holistic approach to security and interacts with other security devices. At the same time it requires significant efforts by the administrator to define and manage all the security policies.

In our implementation the NGFW is directly connected with the Domain Controller to continuously authenticate and authorize the users. After the initial authentication phase managed by the Access Proxy, any time a user wants to interact with another resource in the network the request is directed to the NGFW which inspects the packets. Through the Trust Algorithm, it determines whether that specific user is allowed to perform the request. In such a way, the NGFW can also act as a Policy Administrator with the ability to suspend or interrupt a user session and initiate a new authorization and authentication process in accordance with policy requirements. From the implementation point of view, the role of the NGFW is broader than just monitoring communications between the IT and the OT systems, like in an I-DMZ setup. Indeed, it needs to process all the communications between each device inside the network to implement fine-grained access controls and monitor each communication. The NGFW is physically placed between the IT and the OT to force the communications path, and acts as a router between the different network segments. Essentially, the NGFW operates as a Virtual conduit between the different devices by monitoring every transmitted packet. It can allow or deny traffic based on the zone's VLAN tag to enforce security or dynamically restrict access to a network segment if it is needed. The NGFW leverages IPTables rules for its operations, in particular NFQUEUE an IPTables target that delegates the packet filtering decision to an user-space software written in Python. The packets are blocked in a queue until the Trust Algorithm determines whether the




Fig. 2. A&A Flow

connection can be authorized. After a positive results from all security checks, the software forwards the packets to the original destination while dropping the malicious ones along with a log containing the root cause for the rejection.

C. Trust Algorithm

The Trust Algorithm (TA) is the other fundamental component of our Zero Trust solution. We design and implement it by using two components: the Access Proxy and the Next Generation Firewall. The former uses it for the first authentication and authorization (A&A) phase and determines if a user can interact with the requested service. The latter uses the Trust Algorithm to re-authenticate and re-authorize the connections during a session. As shown in Figure 2, the Policy Engine is the core of the access management. It aggregates all the user, device, and context information to authenticate it and dynamically grants the appropriate authorization level.

The Policy Engine is also the component that executes the Trust Algorithm process that allows or deny the access to a specific resource for every request. This algorithm comprises several steps which must be performed before approving a request. Referring to the NIST definitions of Trust Algorithms, our solution falls into the category of a criteria-based solution with contextual information. Each user has the right to access specific resources, and its behavior is continuously monitored during communications. The TA consists of several steps that can be configured and customized to reflect the policies and procedures defined by the organization. The algorithm is composed by the following steps:

- 1) The user sends a request for a resource.
- 2) The user is authenticated through its credentials.
- 3) The SSO system asks the Domain controller and validates the user credentials.

- 4) If the user is authorized starts the device A&A phase.
- 5) The AP/NGFW starts the operations to identify the device.
- 6) At this point, the AP/NGFW made the authorization check.
- 7) Identify the user's authorization level based on the identity and contextual information.
- 8) Determine the device health and decide if it is sufficiently reliable.
- 9) If all these checks are passed, the AP allows access to the right resource.

The decision to grant access to a resource is based on several contextual data with the threat intelligence systems that play an essential role. This solution collects, aggregates, and organizes threat intelligence data from multiple sources and formats, providing the security teams with information about known malware, IP addresses and other threats. It can execute efficient and accurate threat identification, investigation, and response. This service can help to identify, for example, if a device is affected by a known vulnerability or if is not patched, by dynamically changing the trust level of a specific device. Other possible considered information consists of the user's access history to a resource, such as the time of access, the location and the IP of the request. All these details are stored in a security context to profile the user's behavior, and can generate alerts if something unusual is detected.

The NGFW also offers re-authentication and re-authorization services. The Trust algorithm used by this device is different from that of the Access Proxy because now the goal is to determine whether the communications that pass through the NGFW are still authorized. It checks if the session token released by the Access Proxy has expired or has lost its validity due to a change in the security context. Then it initiates a new authentication and authorization phase and starts to monitor all the following communications between the devices.

In our implementation, the NGFW leverages iptables to block packets in transit until the communication is authenticated and authorized by the policies engine. All the collected information contributes to define the risk level of the connection. The trust score calculated at the end of the algorithm is compared with a dynamic threshold that can change accordingly with the security context of the connection and with the level of privilege required to access a specific resource inside the company's network.

IV. EXPERIMENTAL EVALUATION

This section describes the test performed on the implemented solution that can validate its security properties. We have built a realistic industrial network with multiple devices, including a compromised one, to see whether the countermeasures are effective in detecting and reacting to cyber threats. As described in Section III, we relied on two main security components to implement the Zero Trust core principles: an Access Proxy and a Next Generation Firewall. For the AP, we have used Flask [9] to create a Single-Sign-On system with a

login page and a back-end that runs the Trust Algorithm for the authorization and authentication process. The NGFW test environment comprises a PLC and a compromised IT machine (Ubuntu Virtual machine) placed in two different network segments, with another Ubuntu Virtual machine that acts as NGFW.

a) Access Proxy: To test our Access proxy, we employed a proof of concept of a Credential-Based Attack. We assume that a malicious actor managed to steal a set of credentials and tries to use them to gain access to the organization's network. Once the user enters the credentials, the system sends a query to the domain controller to verify the user's permissions inside the system. Then, we identify the device and all the other relevant information about the environment. We use Nmap with the "-O" flag, which stands for "OS detection" [20] to retrieve the CPE of the device and other insights that can be used for multiple security purposes. We can use those data to determine the vulnerabilities of target hosts, and by analyzing the asset database, we can detect unauthorized and dangerous devices, which is our primary goal. All those data, together with the threat intelligence systems, allow us to define the context of the communications better, determine the health of the device, and block ambiguous connection attempts. For this proof of concept, we have simulated the threat intelligence system using two databases. The first database is provided and maintained by the Cybersecurity & Infrastructure Security Agency (CISA) and is the "*Known Exploited Vulnerabilities Catalog*" [4] used to extract all the vulnerabilities discovered so far. The second database is the "*National Vulnerability Database*" that is provided by the NIST. The agency exposes an API [19] through which we can retrieve all the information regarding a specific CVE to check the vulnerabilities and the associated risks.

The API returns helpful information that can be used to set more fine-grained controls according to company policies. In our case, we save the baseScore, the exploitabilityScore, the impactScore, and the AttackVector. Even if the AttackVector value is included in the calculation of the final impact score, we want to give more emphasis to this value. This is because when evaluating the risk associated with remote access, this value reflects the context by which vulnerability exploitation is possible; in particular, this value will be greater the higher the risk of an exploit by a remote attacker. Finally, we compare the calculated trust score with a threshold that can be dynamically modified according to the context and resource accessed. At this point, we use the JWT module to generate the encrypted token, which will be associated with the user's traffic to trace all the operations. The Trust Algorithm approach's flexibility allows us to increase or decrease the granularity of our controls according to the context of the specific environment.

b) NGFW: To test the behavior of the Next Generation Firewall, we assume the possibility of sending Modbus commands from a machine in the IT department to a PLC inside the OT network. Therefore the NGFW should understand not only if the user is logged into the IT machine but also if it is authorized to perform this type of action and if the

machine from which the connection comes is sufficiently reliable. This is crucial since, looking at the Tactics, Techniques, and Procedures (TTPs), IT machines are the first to be compromised during an attack on an industrial plant. To validate our proposal, we assume that a malicious actor has managed to take control of a machine in the IT network, so we are in an advanced step of a typical industrial cyber attack. As previously said, we leverage IpTables functionalities to block the in-transit packets until the Trust algorithm is completed and a verdict can be emitted. Once we have captured the packet in the queue, thanks to NETFILTER hooks, we can manage it. The Netfilter hooks [18] is a framework inside the Linux kernel that allows kernel modules to register callback functions at different locations of the Linux network stack. The registered callback function is then called back for every packet that traverses the respective hook within the Linux network stack. Netfilter provides access to packets matched by an IPTables rule in Linux. In this way, matched packets can be accepted, dropped, altered, reordered or marked for further inspection. The system then determines which employee is logged in to the machine and queries the domain controller to retrieve the user's access privilege for the underlying resource. At this step, it is also possible to set additional environment information and define more fine-grained policies to identify an employee. An example is the ability to check the connection's time of the day and compare it with the employee's working hours to detect anomalies.

Once these checks are passed, we verify the device's health with the same algorithm used by the Access Proxy IV-0a and we emit a positive or negative verdict to authorize or not the connection. Figure 3 shows a log of the policies engine referring to a connection attempt through a Windows 7 machine, clearly affected by many vulnerabilities, which will not be allowed to access due to a too-low trust level. The result of a blocked connection attempt is shown in Figure 4, the NGFW was able to detect the threat and block it. In this case, we can see that the IT machine sends a lot of TCP re-transmissions denoting that the firewall has done its job.

A. Comparison with I-DMZ

The results of our tests showed the advantages of the Zero Trust approach with respect to the I-DMZ solution. Organizations can exploit our approach based on a modular architecture to securely support their digital transformations without restrictions on new services, which can be implemented without any compromise in terms of security. The combined work of Access Proxy and NGFW allows fine-grained and flexible internal data flow management without any architectural restrictions. A standard I-DMZ method which bases its security model on a rigid layered architecture is no longer suitable for the modern industrial sector. With the adoption of cloud, edge computing and increasingly connected services, vast amounts of data can also be collected at Level 1, processed, and sent directly to the cloud bypassing the hierarchical data flows of the Purdue model. The assumption underlying the I-DMZ approach that only the IT part of

```

this ip address is associated with silviorusso
the work hour is correct

Name: Microsoft Windows Group Policy Privilege Escalation
CVE: CVE-2014-1812
Base score: 9.0
exploitability score: 8.0
impact score: 10.0
attack vector NETWORK
your device version is patched for the vulnerability: CVE-2014-1812

Name: Microsoft Windows Kernel 'Win32k.sys' Local Privilege Escalation Vulnerability
CVE: CVE-2016-0167
Base score: 7.2
exploitability score: 3.9
impact score: 10.0
attack vector LOCAL
your device version is patched for the vulnerability: CVE-2016-0167

Name: Microsoft Windows Kernel Information Disclosure Vulnerability
CVE: CVE-2021-31955
Base score: 2.1
exploitability score: 3.9
impact score: 2.9
attack vector LOCAL
your device version is patched for the vulnerability: CVE-2021-31955

Name: Microsoft Windows Media Center Remote Code Execution vulnerability
CVE: CVE-2016-0185
Base score: 9.3
exploitability score: 8.6
impact score: 10.0
attack vector NETWORK
your device version is not patched for the vulnerability: CVE-2016-0185

Name: Microsoft Windows Installer Privilege Escalation Vulnerability
CVE: CVE-2020-0683
Base score: 7.2
exploitability score: 3.9
impact score: 10.0
attack vector LOCAL
your device version is patched for the vulnerability: CVE-2020-0683
The final trust level is -9.84

```

Fig. 3. Windows machine connection attempt

No.	Time	Source	Destination	Prot	Ler	Info
5	3.175730717	172.254.179.2	169.254.179.67	TCP	74	47851 → 502 [SYN] Seq
6	4.196436181	172.254.179.2	169.254.179.67	TCP	74	[TCP Retransmission]
13	6.179329398	172.254.179.2	169.254.179.67	TCP	74	56933 → 502 [SYN] Seq
14	7.204542982	172.254.179.2	169.254.179.67	TCP	74	[TCP Retransmission]
26	22.840859816	172.254.179.2	169.254.179.67	TCP	74	33143 → 502 [SYN] Seq
27	23.844478430	172.254.179.2	169.254.179.67	TCP	74	[TCP Retransmission]
28	25.843583960	172.254.179.2	169.254.179.67	TCP	74	53167 → 502 [SYN] Seq
31	26.852491574	172.254.179.2	169.254.179.67	TCP	74	[TCP Retransmission]

Fig. 4. TCP connection blocked from the Firewall.

the architecture can be connected with the external world is broken. Therefore, this model can no longer be used to effectively protect our architectures without imposing hard limits on the digital development of organizations. Comparing I-DMZ with our approach, the paradigm shift is evident. While in the former, it is the data flow that must be adapted to the architecture in order to guarantee its security, in our case it is the architecture itself that adapts its behavior to the needs of the organization thanks to the policy engine and the trust algorithm.

Another significant benefit of our solution is the simplification of the hardening process. One of the most critical problems for the OT sector is securing legacy systems and keeping them updated. With the I-DMZ approach, the hardening process requires an architectural change that is not always easy to implement, leaving these systems without the necessary protection. Our solution and the ZT approach in general allow for dynamically adjusting the security controls on a per-resource basis to match its risk level.

Our experiments highlight the necessity even for the indus-

trial sector to move from a rigid layered architecture based on the Purdue model to a more flexible approach to cybersecurity. Otherwise, organizations will be forced to choose between the security of their infrastructure and the ability to innovate by integrating bleeding-edge technologies.

V. CONCLUSIONS

The proposed Zero Trust solution defines a novel and innovative way to structure an industrial network. We have actively removed the concept of internal perimeter to rely on a new approach identity-centric where there are no implicit trusted zones and everything must be authenticated, authorized and monitored. The rapid adoption of emerging technologies to facilitate remote connectivity even in industrial contexts (IIoT) force the organization to review and update existing security strategies and countermeasures to protect the most critical systems. Standard defense practices, such as industrial DMZ (I-DMZ), are insufficient for guaranteeing security and flexibility. IIoT solutions and the rise of remote work are blurring the separation between the trusted internal network and the public internet, which is a core assumption for the I-DMZ based security model.

Our solution is inspired by state-of-the-art best practices of Zero Trust Architectures for the IT sector. Moreover, it is specifically customized for OT systems to bring all the benefits of this new paradigm to the industrial sector. The achieved result can be considered a starting point from which any company can redesign its security infrastructure. Different from other approaches, the Zero Trust security model does not impose rigid constraints on the architecture; it is a moving target for the security posture of an entire organization. The implementation of a Zero Trust Architecture should be tailored for the specific context by carefully prioritizing the appropriate security tenets in order to minimize the overall risk. Our research evidenced that by embracing the Zero Trust approach industries are forced to adopt all the best practices for a well-defined cybersecurity strategy. During the implementation and validation phases, it has become clear that for the transition to a Zero Trust approach, the economic factor is not the main obstacle to its adoption [1]. The real limiting factor is the total effort needed both in time and skill to implement it. The Zero Trust model cannot be treated as a black box that can be turned on at will, all the choices must be correctly planned and examined, in addition a deep knowledge of the underlying infrastructure and all the interactions between the different components is required for a successful implementation.

REFERENCES

- [1] Z. Adahman, A. W. Malik, and Z. Anwar, "An analysis of zero-trust architecture and its cost-effectiveness for organizational security," *Computers Security*, vol. 122, p. 102911, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167404822003042>
- [2] A. Alagappan, S. K. Venkatachary, and L. J. B. Andrews, "Augmenting zero trust network architecture to enhance security in virtual power plants," *Energy Reports*, vol. 8, pp. 1309–1320, 2022.
- [3] J. R. N. Alvarez, Y. P. Zamora, I. B. Pina, and E. N. Angarita, "Demilitarized network to secure the data stored in industrial networks," *International Journal of Electrical & Computer Engineering (2088-8708)*, vol. 11, no. 1, 2021.
- [4] CISA. Known exploited vulnerabilities catalog. [Online]. Available: <https://www.cisa.gov/known-exploited-vulnerabilities-catalog>
- [5] G. W. S. E. Dan Perez, Sarah Jones. Check your pulse: Suspected apt actors leverage authentication bypass techniques and pulse secure zero-day. [Online]. Available: <https://www.mandiant.com/resources/blog/suspected-apt-actors-leverage-bypass-techniques-pulse-secure-zero-day>
- [6] G. W. S. E. E. H. Dan Perez, Sarah Jones. Re-checking your pulse: Updates on chinese apt actors compromising pulse secure vpn devices. [Online]. Available: <https://www.mandiant.com/resources/blog/updates-on-chinese-apt-compromising-pulse-secure-vpn-devices>
- [7] S. D. Y. A. Director. (2022) Moving the u.s. government toward zero trust cybersecurity principles. [Online]. Available: <https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf>
- [8] W. Fang and X. Guan, "Research on ios remote security access technology based on zero trust," in *2022 IEEE 6th Information Technology and Mechatronics Engineering Conference (ITOEC)*, vol. 6, March 2022, pp. 238–241.
- [9] Flask. (2022) Flask project. [Online]. Available: <https://flask.palletsprojects.com/en/2.2.x>
- [10] "IEC/TS 62443 Industrial communication networks – Network and system security," INTERNATIONAL ELECTROTECHNICAL COMMISSION, Standard.
- [11] N. Jiang, H. Lin, Z. Yin, and C. Xi, "Research of paired industrial firewalls in defense-in-depth architecture of integrated manufacturing or production system," in *2017 IEEE International Conference on Information and Automation (ICIA)*, 2017, pp. 523–526.
- [12] N. Jiang, H. Lin, Z. Yin, and L. Zheng, "Performance research on industrial demilitarized zone in defense-in-depth architecture," in *2018 IEEE International Conference on Information and Automation (ICIA)*, 2018, pp. 534–537.
- [13] P. Katuruza. It of convergence, managing the cybersecurity risks. [Online]. Available: <https://gca.isa.org/blog/it-ot-convergence-managing-the-cybersecurity-risks>
- [14] J. Kindervag *et al.*, "No more chewy centers: The zero trust model of information security," *Forrester Research, Inc., dated Mar.*, vol. 23, 2016.
- [15] F. Magnanini, L. Ferretti, and M. Colajanni, "Flexible and survivable single sign-on," in *Cyberspace Safety and Security*, W. Meng and M. Conti, Eds. Cham: Springer International Publishing, 2022, pp. 182–197.
- [16] A. W. Mir and K. R. Ram Kumar, "Zero trust user access and identity security in smart grid based scada systems," in *Proceedings of the 12th International Conference on Soft Computing and Pattern Recognition (SoCPar 2020)*, A. Abraham, Y. Ohswawa, N. Gandhi, M. Jabbar, A. Haqiq, S. McLoone, and B. Issac, Eds. Cham: Springer International Publishing, 2021, pp. 716–726.
- [17] G. Murray, M. N. Johnstone, and C. Valli, "The convergence of IT and OT in critical infrastructure." Security Research Institute (SRI), Edith Cowan University, 2017.
- [18] netfilter.org. Iptables. [Online]. Available: <https://www.netfilter.org>
- [19] NIST. National vulnerability database. [Online]. Available: <https://nvd.nist.gov/developers/vulnerabilities>
- [20] Nmap.org. Remote os detection. [Online]. Available: <https://nmap.org/book/osdetect.html#idm45323735721296>
- [21] D. Puthal, L. T. Yang, S. Dustdar, Z. Wen, S. Jun, A. v. Moorsel, and R. Ranjan, "A user-centric security solution for internet of things and edge convergence," *ACM Trans. Cyber-Phys. Syst.*, vol. 4, no. 3, may 2020. [Online]. Available: <https://doi.org/10.1145/3351882>
- [22] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero trust architecture," 2020-08-10 04:08:00 2020.
- [23] N. F. Syed, S. W. Shah, A. Shaghghi, A. Anwar, Z. Baig, and R. Doss, "Zero trust architecture (zta): A comprehensive survey," *IEEE Access*, vol. 10, pp. 57 143–57 179, 2022.
- [24] R. L. Tyler Mclellan, Justin Moore. Unc2447 sombrat and fivehands ransomware: A sophisticated financial threat. [Online]. Available: <https://www.mandiant.com/resources/blog/unc2447-sombrat-and-fivehands-ransomware-sophisticated-financial-threat>
- [25] R. Ward and B. Beyer, "Beyondcorp: A new approach to enterprise security," vol. Vol. 39, No. 6, pp. 6–11, 2014.
- [26] Zscaler. (2021) Vpn risk report. [Online]. Available: <https://info.zscaler.com/rs/306-ZEJ-256/images/2021-VPN-Risk-Report-Zscaler.pdf>