Technical Disclosure Commons

Defensive Publications Series

June 2022

NETWORK-EMBEDDED ZERO TRUST AGENT FOR AUTOMATION AND CONTROL SYSTEMS

Maik Seewald

Laurent Hausermann

Andre Guerard

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Recommended Citation
Seewald, Maik; Hausermann, Laurent; and Guerard, Andre, "NETWORK-EMBEDDED ZERO TRUST AGENT FOR AUTOMATION AND CONTROL SYSTEMS", Technical Disclosure Commons, (June 13, 2022)
https://www.tdcommons.org/dpubs_series/5196

This work is licensed under a Creative Commons Attribution 4.0 License.
This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

NETWORK-EMBEDDED ZERO TRUST AGENT FOR AUTOMATION AND CONTROL SYSTEMS

AUTHORS:
Maik Seewald
Laurent Hausermann
Andre Guerard

ABSTRACT

Techniques are presented herein that support a zero trust architecture for (e.g., legacy) devices in an operational technology (OT) and industrial internet of things (IIoT) environment. The presented techniques support granular segmentation and access control to critical assets based on a network equipment vendor's ability to host applications (in this instance a proxy) on network components such as switches. Aspects of the presented techniques comprise three key components or artifacts. A first artifact encompasses a proxy component that is hosted as part of a control system. A second artifact encompasses a proxy component – i.e., a zero trust agent (ZTA) – that is hosted on a switch, and which enforces the (e.g., role-based access control (RBAC)) security of a cabinet which will be considered a resource enclave, and which may be assigned a cryptographically-based identity. A third artifact encompasses a policy administration and authentication server.

DETAILED DESCRIPTION

Current trends towards the convergence of information technology (IT) and operational technology (OT), in addition to a growing threat environment, require very granular network segmentation and access control to devices, applications, and data. A traditional segmentation model that is based on zones and access control is not sufficient to reach the security level that is necessary. This is especially important for critical processes in industrial and power utility automation and control. Devices (e.g., programmable logic controllers (PLCs), remote terminal units (RTUs), sensors, etc.) in an industrial environment typically have a lifetime that is greater than 10 or even 20 years, often with little implemented security, limited computing resources, and no ability to receive software updates to meet access control and segmentation needs. Moreover, the

creation of a secure ecosystem of devices in OT systems faces a chicken and egg problem – i.e., original equipment manufacturers (OEMs) are waiting for customer requests before implementing a solution, and customers are waiting for a complete solution that is ready to deploy. In the end, even if a fully secured solution were to exist the inertia of such systems dictates that existing unsecured devices, often called legacy devices, will stay in operation for the next 10 to 20 years, depending upon the involved business segment. For the most critical industries (such as, for example, nuclear power) such a timeframe may be 40 years. Therefore, a solution is needed for establishing a zero trust architecture in an industrial system in a less invasive way than redesigning all the OT components.

To address such a challenge, techniques are presented herein that enable a zero trust architecture for (e.g., legacy) devices in an OT and industrial internet of things (IIoT) environment. Aspects of the presented techniques support a scalable security architecture as well as the protocols that are typically used in such industries.

Regarding the current state of the art in industrial security, the National Institute of Standards and Technology (NIST) special publication (SP) 800-27 identifies a number of zero trust tenets. Those tenets include:

1. All data sources and computing services are considered resources.
2. All communication is secured regardless of network location.
3. Access to individual enterprise resources is granted on a per-session basis.
4. Access to resources is determined by a dynamic policy.
5. An enterprise monitors and measures the integrity and security posture of all of the owned and associated assets.
6. All resource authentication and authorization are dynamic and are strictly enforced before access is allowed.
7. An enterprise collects as much information as possible about the current state of assets, the network infrastructure, and communications and uses that information to improve its security posture.

As described in the introductory chapters of NIST SP 800-27, the zero trust objective encompasses "prevent[ing] unauthorized access to data and services coupled with making the access control enforcement as granular as possible." Figure 1, below, presents elements of an illustrative system that is reflective of the above discussion.

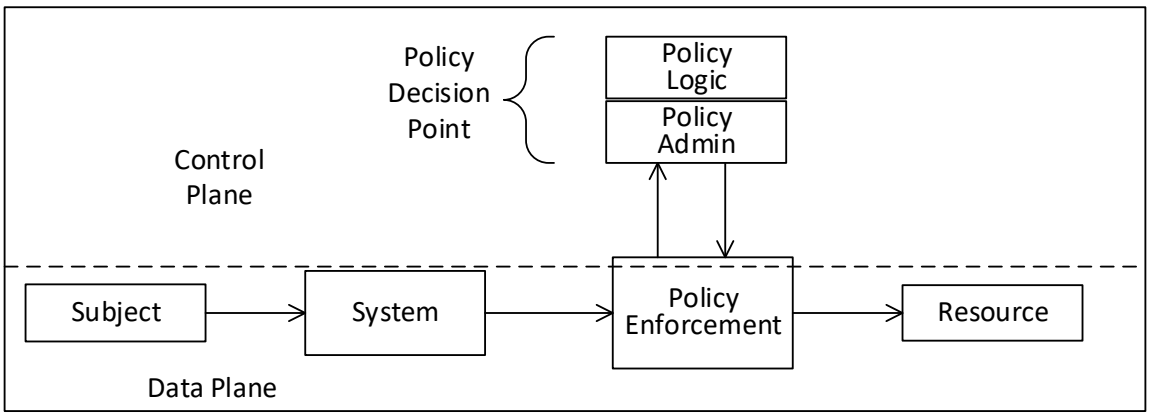

Figure 1: Illustrative Zero Trust System

As depicted in Figure 1, above, a subject that wishes to gain access to an enterprise resource proceeds through a Policy Enforcement Point (PEP) that either allows or disallows access to the resource (based on its identity, role, or other factors) as computed by a Policy Decision Point (PDP).

The above concepts are not new and have been theorized by the Organization for the Advancement of Structured Information Standards (OASIS) eXtensible Access Control Markup Language (XACML) standard. For example, in connection with energy automation the concepts are specified in Part 8 of the International Electrotechnical Commission (IEC) security standard 62351 (i.e., 62351-8). According to that standard, each device has a PEP and PDP engine and is able to allow access to its resources based on the identity and role of the peer that is trying to connect. It is important to note that for the model to be complete, not only human actors, but also any devices that are able to initiate a connection, must have an identity. In 62351-8, such entities are summarized by the name "Subject." A resource's access may be defined by not only the role of the user but also by the status of the device on which they apply (e.g., see IEC 62351-90-1 for information on role assignments).

Thus, the IEC security standard 62351-8 defines a zero trust architecture to which the NIST has added monitoring concepts. Each intelligent electronic device (IED), such as a digital protection relay or other control device, is considered to be a resource that both a

zero trust model and the IEC security standard 62351 may rely upon according to the theory that is established by the OASIS XACML standard.

It is important to note at this point that the type of fine granular role-based access control (as described above) cannot be replaced by a simple authentication (when a user or a device connects to a network) comprising a fixed network policy through routing or an access-control list (ACL). Such policies are static, not fine-grained, and do not comply with the concept of session and security of communications between peers.

For existing legacy systems, the NIST publications propose a not-so-ideal solution. That solution, the Enclave Gateway Model, is presented in Figure 2, below.

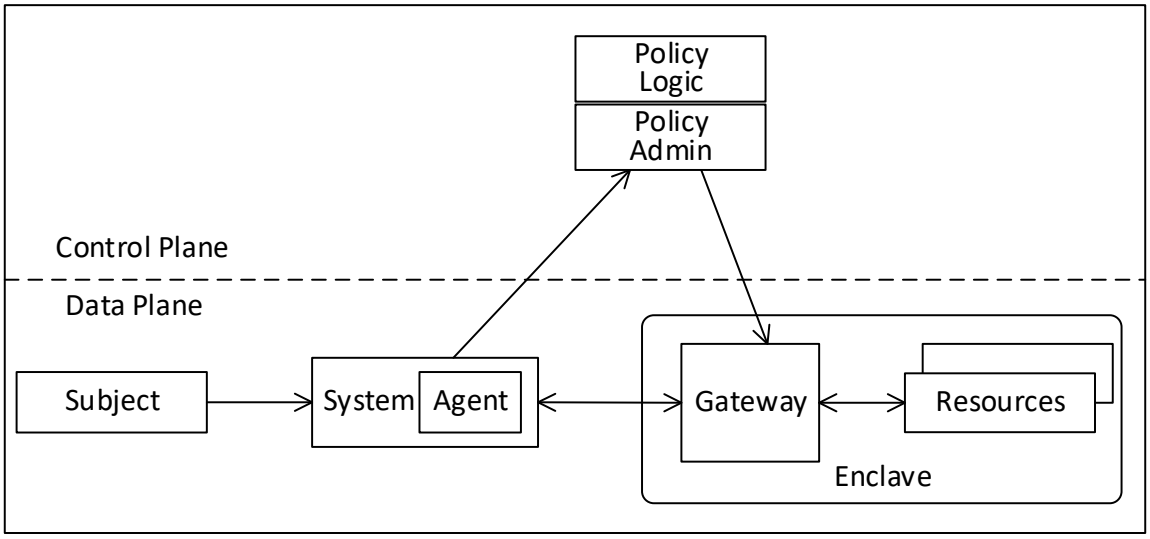

Figure 2: Enclave Gateway Model

Under the Enclave Gateway Model, as depicted in Figure 2, above, access to the different resources is protected through a zero trust agent or gateway pair that provides identity checks, session establishment, and RBAC enforcement. However, a problem arises regarding granularity and the number of resources. Figure 3, below, depicts elements of an architecture for a "classical industrial system."

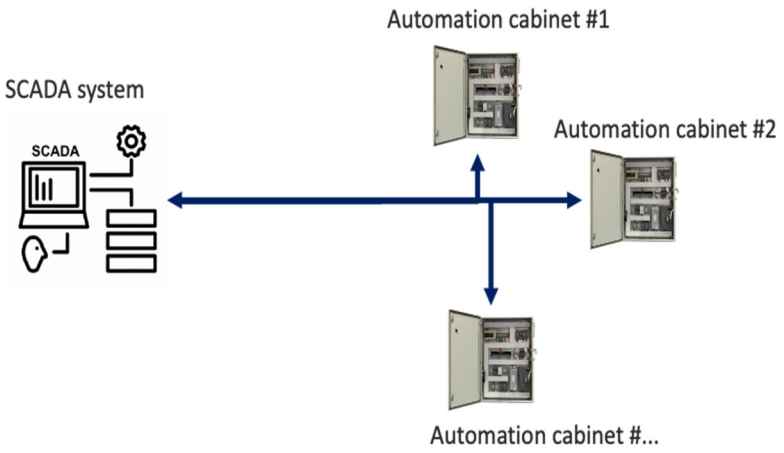

Figure 3: Classical Industrial System Architecture

As illustrated in Figure 3, above, control systems (such as an energy management system (EMS), a supervisory control and data acquisition (SCADA) system, etc.) are typically host-based applications that are connected to a series of automation cabinets that are fixed (and theoretically physically protected) subsystems. Each cabinet is usually assigned to a specific functional task for automation and control. It is important to note that each of the automation cabinets may contain one or more industrial switches to enable connectivity between multiple automation and control devices (e.g., PLCs, IEDs, RTUs) and the central control system.

Aspects of the techniques presented herein, which will be described and illustrated in the below narrative, comprise three key components or features. A first feature encompasses a proxy component that is hosted as part of a control system. A second feature encompasses a proxy component – i.e., a zero trust agent (ZTA) – that is hosted on a switch, and which enforces the security of a cabinet which will be considered a resource enclave. This comprises granular RBAC for connected automation and control devices as well as integrated support for physical security. Moreover, the ZTA may provide a cryptographically based identity to each secured cabinet, providing for a fine grained zero trust architecture. A third feature encompasses a policy administration and authentication server.

According to aspects of the techniques presented herein, a ZTA may be hosted on a network equipment vendor's application environment and may support multi-layer industrial network security by incorporating zero trust network (ZTN) capabilities in

addition to traditional network security (e.g., ACLs). Such a ZTA offers a number of functionalities.

Under a first functionality, a ZTA enables session-based transport security, between a cabinet and the control system, which is based on Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) in combination with the proxy component running in the context of the control system.

Under a next functionality, a ZTA performs the security policy enforcement of incoming connections. In more detail, the ZTA supports the enforcement of RBAC for each protocol (such as File Transfer Protocol (FTP), the World Wide Web, SCADA, IEC Fieldbus, IEC 61850, etc.) and for each addressed device (i.e., resource) with the help of the identity of a requester and a submitted credential (e.g., a certificate or an attribute certificate). Such enforcement is described below.

First, as noted previously, RBAC is a mechanism that enables fine granular access control based on roles that are associated to resources. Typical roles include administrator, operator, or installer (e.g., a technician).

Second, RBAC for power utility automation is defined in IEC 62351-8 where the specification presupposes an implementation in the device or resource (typically in an IED or an RTU). The majority of the devices do not yet implement RBAC and, accordingly, they are not capable of receiving software or firmware updates. In this sense, aspects of the techniques presented herein provide a way for retrofitting a large number of devices since an upgrade to RBAC on the host (i.e., a control system) will typically comprise just a software update.

Third, in order to enforce RBAC a ZTA implements two core functions. A first core function encompasses receiving a credential (e.g., an attribute certificate or a token) and encapsulated role names as part of a TLS or DTLS transmission. A second core function encompasses authorizing and allowing access to device resources based on the roles, where a roles-to-permission mapping (including the data access) is based on the device data model which is part of the ZTA configuration. Figure 4, below, presents elements of a functional architecture as described above.

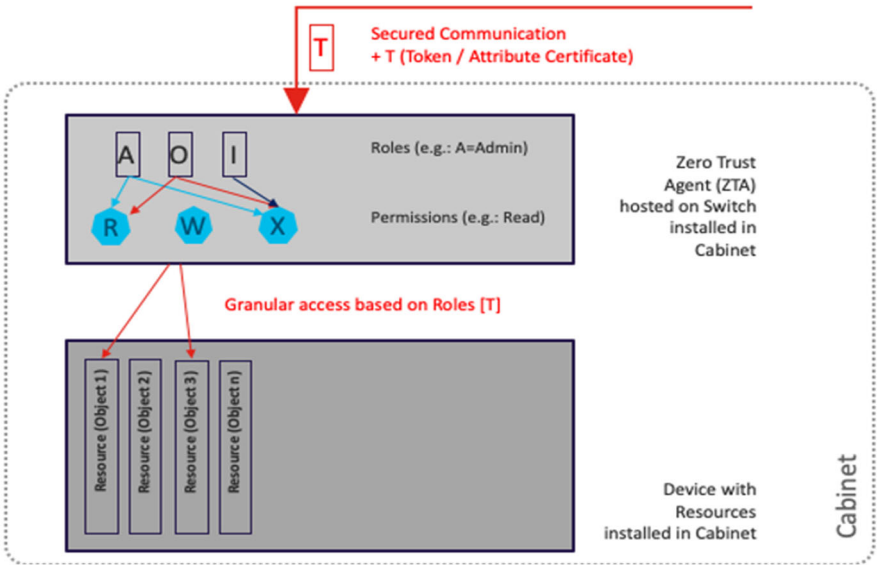

Figure 4: Illustrative Functional Architecture

Fourth, during operation a ZTA executes RBAC based on the role (which is part of, or in the content of, a credential). While RBAC in an end device (e.g., an IED) is known, a proxy-based implementation (according to aspects of the techniques presented herein and as described and illustrated above) is innovative.

Fifth, RBAC in an entire system is only possible if all of the components of the system have both an identity and an associated role. By providing such an identity to each automation cabinet (including the devices) and associating roles through a policy administration server, a ZTA allows for a real-world implementation of zero trust concepts in existing automation systems. Each element of such a system is authorized to perform only that for which it has been designed. An identity, as described above, that is given to each control cabinet and each device within a cabinet is innovative.

Under a further functionality, in order to complete the risk model, the physical security of a cabinet is also enforced. Physical accesses in an automation system are often ignored but are one of the major threats and sources of failures. By combining an RBAC, a global security model, and a physical security mechanism, a ZTA brings the concept of zero trust into the physical world and into the daily challenges of industrial systems. Such cyber-physical control of an automation cabinet is innovative and the monitoring aspect of same may be ensured through, for example, a sensor in a switch.

According to aspects of the techniques presented herein, a ZTA (proxy component) may have a software pendant for host-based applications. A host (e.g., an industrial personal computer (PC)) itself may be considered as a resource enclave. The identity that is used may be the identity of the proxy (i.e., authenticating the machine) or the identity of the human that is using it (i.e., an operator).

Under further aspects of the techniques presented herein, all of the proxies may be centrally managed by a policy administration and authentication server which may be an identity services engine function or component. Such a server may provide an identity and role to all of the proxies and users of a system through classic cryptographic mechanisms, provide a configuration mechanism for RBAC in all of the proxies, and provide the credentials (e.g., attribute certificates and tokens) which contain the role name itself.

Figure 5, below, depicts elements of a deployment example in an industrial automation setting according to aspects of the techniques presented herein and reflective of the above discussion.

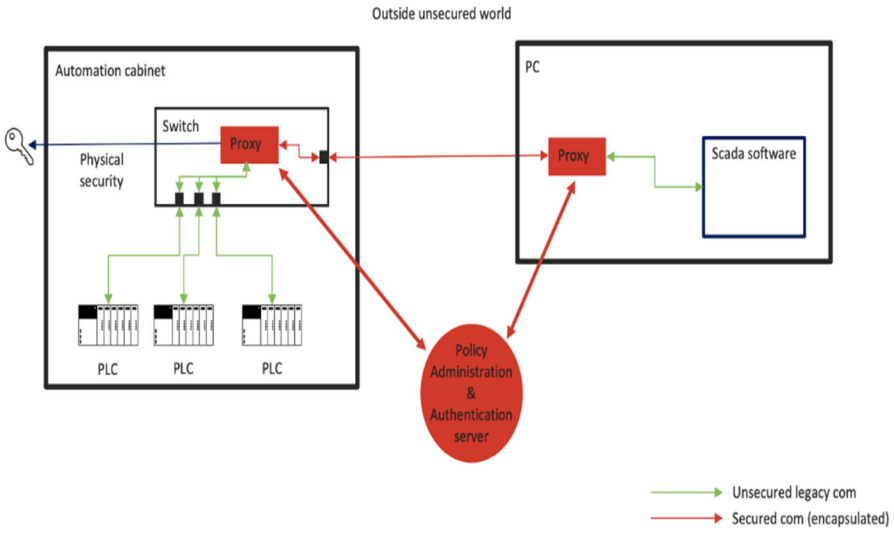

Figure 5: Exemplary Industrial Automation Deployment

The exemplary deployment that was illustrated in Figure 5, above, could be applied as well to a power utility automation context (connecting through a proxy to digital protection and control devices). In such a case, a cabinet could be a bay (i.e., part of an

electrical substation). Through such means, an asset owner and operator may (by adding, using, and configuring ZTAs on switches) establish a zero trust architecture in the control system without having to rebuild and change all of the components and devices of an automation system. Additionally, aspects of the techniques presented herein allow for a flexible, policy-based configuration of different security levels (SLs) for each cabinet. Such an arrangement supports strong security in critical subparts of a system while making other parts of the system more easily accessible.

In summary, techniques have been presented herein that support zero trust architecture for (e.g., legacy) devices in an OT and IIoT environment. The presented techniques support granular segmentation and access control to critical assets based on a network equipment vendor's ability to host applications (in this instance a proxy) on network components such as switches and routers. Aspects of the presented techniques comprise three key components or artifacts. A first artifact encompasses a proxy component that is hosted as part of a control system. A second artifact encompasses a proxy component – i.e., a ZTA – that is hosted on a switch, and which enforces the (e.g., RBAC) security of a cabinet which will be considered a resource enclave, and which may be assigned a cryptographically-based identity. A third artifact encompasses a policy administration and authentication server.