Received 2 May 2023; revised 8 August 2023 and 21 August 2023; accepted 21 August 2023. Date of publication 29 August 2023; date of current version 19 September 2023. The review of this article was arranged by Associate Editor P. Zanchetta.

Digital Object Identifier 10.1109/OJIA.2023.3309669

Evolution and Trends of Cloud on Industrial OT Networks

CYRIL PERDUCAT [1] (Member, IEEE), DAVID C. MAZUR [1] (Senior Member, IEEE), WES MUKAI [2] (Member, IEEE), SCOTT N. SANDLER [3] (Senior Member, IEEE), MICHAEL J. ANTHONY [1] (Member, IEEE), AND JON A. MILLS [3] (Senior Member, IEEE)

[1]Rockwell Automation, Milwaukee, WI 53204 USA
[2]Rockwell Automation, San Jose, CA 95113 USA
[3]Rockwell Automation, Mayfield Heights, OH 44124 USA

CORRESPONDING AUTHOR: DAVID C. MAZUR (e-mail: dcmazur@rockwellautomation.com)

ABSTRACT The current industrial automation landscape faces considerable challenges due to the increasing growth of Industrial Internet of Things devices, cloud services, information technology (IT)/operational technology (OT) convergence, along with evolving hyperscaler technologies, such as Kubernetes and distributed computing. This article provides an in-depth review of the existing Instrumentation Society of America (ISA)-95 Model and its current role in supporting manufacturing systems and their interactions. Additionally, it examines how emerging technologies are impacting the security, design, and management of OT networks. As existing perimeter-based models, such as ISA-95, are pushed to their limits, the concepts of zero-trust architectures, and policy-based or software-defined networks are being explored as the next generation of OT network design. This article aims to provide a high-level introduction to the concepts and disruptive technologies, and introduce the potential implementation risks and challenges of these principles within traditional IT/OT converged solutions.

INDEX TERMS Cloud convergence, distributed computing, industrial operational technology (OT) networks, Instrumentation Society of America (ISA)-95, information technology (IT)/OT convergence, Kubernetes, zero trust (ZT).

I. INTRODUCTION

The success or failure of any industrial automation system process is a combination of the electrical and mechanical subprocesses and the human input working together effectively. This coordination historically had been the responsibility of a joint effort between engineering, operations, and maintenance working together to define the correct outcomes for the process. As production evolved from hardwired instrument panels to fieldbus, additional considerations need to be provided to ensure proper communication between systems running both proprietary and open protocols.

With the introduction of industrial ethernet, larger amounts of devices were introduced into system architectures. This increased volume of ethernet devices created a need for standards and guidelines to ensure the stability and throughput of networks. The most prevalent of these standards and guidance came from the Instrumentation Society of America (ISA) ISA-95.

Intelligent devices within industrial automation systems have become more data rich. Subsequently the needs of applications, analytic systems, and data platforms have become more demanding for data flow both on premise, at the edge, and within cloud environments. As edge and cloud architectures become more prevalent in industrial automation, there are new considerations on network performance, security, and resiliency that need to be considered.

With this increased need for data from OT-based systems, in conjunction with the evolving workforce challenges and shortages, remote access and multiple networks have become a reality. This evolution has pushed current security practices to their limits, thus introducing new concepts, such as policy-based software defined networks, zero trust (ZT), and workload authorization [1], [2].

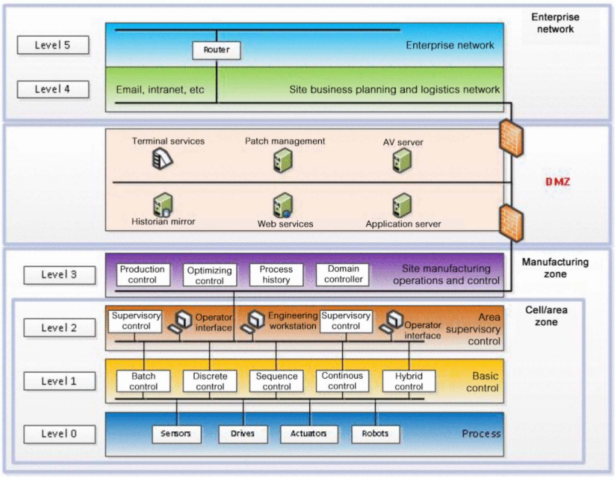

FIGURE 1. **ISA-95 model levels.**

This article intends to provide an overview of the current strategies in place within OT networks, and provide an introduction to the new and disruptive technologies that are challenging those strategies, as well as patterns and practices intended to address those challenges.

II. OVERVIEW OF ISA-95/PURDUE MODEL

The ISA-95 Model is a hierarchical model for representing manufacturing systems and their interactions with the enterprise. It is based on the ISA-95 standard (also known as the ANSI/ISA-95 standard), which is an international standard for the integration of enterprise and control systems in the manufacturing and process industries [3]. The standard working group published the first release of ISA-95 in 1995 as "Enterprise Control System Integration" [4].

The goal of the organization was to standardize communications of systems and subsystems in an enterprise-wide environment [5]. Creating this standardization would enable predictable and scalable behaviors of automation control systems ensuring system designers would have reliable rules and best practices to design against for a successful automation system implementation.

The ISA-95 Model is used to represent the relationships between different systems and processes in a manufacturing environment, and to facilitate the integration of these systems and processes. It helps to ensure that the various systems and processes in a manufacturing facility are working together in a coordinated and efficient manner, and that the information flows between these systems and processes are properly managed.

A. LAYERS OF THE ISA-95 MODEL

The ISA-95 Model as seen in Fig. 1, consists of six levels, each representing a different aspect of the manufacturing process:

1) Level 0—Process Level.
2) Level 1—Basic Control Level.
3) Level 2—Supervisory Control Level.
4) Level 3.5—Industrial Demilitarized Zone.
5) Level 4—Site Business and Logistics.
6) Level 5—Enterprise Network.

The lowest level of the ISA-95 architecture is the process level. Components in this level are typically closest to the actual mechanical process. The speed in which the devices act is quasi real time and physically drive a process input or output. Examples of devices within this are sensors, actuators, drives, robots, etc.

Level one of the ISA-95 model defines the basic control zone. Control is typically applied with either a hardware controller, such as a programmable logic controller (PLC), or a soft controller, such as a container or computer implementation of an IEC 61131 controller. Various types of control are housed there including batch, discrete, sequence, continuous, and hybrid control schemas taking input from the lower layer, running logic, and providing outputs to level zero devices.

Level two of the model is known as the supervisory control zone. At this level, operators and engineers have a day-to-day interface with the system. Operator interfaces or human–machine interfaces (HMI) traditionally reside in this area. Furthermore, engineering and operator workstations serve as portals into the process. For larger systems, supervisory controllers may exist to coordinate multiple controllers at level one.

Level three is known as the site manufacturing operations and control zone. Traditional on-premises operational technology (OT) software is often housed at this layer. This can include local manufacturing execution system servers, optimization engines, relational persistence databases, time-series historians, and the domain controller for the OT network.

Level 3.5 is known as the industrial demilitarized zone. This level is often referenced as the line of demarcation between information technology (IT) and OT responsibilities within organizations. The industrial demilitarized zone (IDMZ) is also the barrier of the "trusted" OT network layers below and the "untrusted" site IT network layers above [6].

Level four is commonly referred to as the site business planning and logistics zone. For most organizations, this is often referred to as the intranet and house IT services, such as email, authorization, and access.

Finally, level five is known as the enterprise network and has connectivity to the wide area network. This is often the interface with the internet and other multisite systems within an enterprise.

Overall, the Purdue Model and the ISA-95 standard are important tools for improving the efficiency, reliability, and competitiveness of manufacturing systems. They help to ensure that manufacturing facilities can operate at their full potential and produce high-quality products in a cost-effective manner.

B. SIGNIFICANCE OF THE INDUSTRIAL DEMILITARIZED ZONE

The IDMZ is used as a barrier between the enterprise and industrial zones. This concept is commonplace in traditional IT networks but is newer within the industrial automation community. With cyber security attacks at an all-time high worldwide, special concern needs to be placed in how to best secure industrial processes [7]. It should be noted that IDMZs are a form of physical access control, not application control [6].

The demilitarized zone plays a significant role in the separation of IT and OT responsibility within an industrial automation system. Standard practice had been to design the OT network, levels 0–3 as trusted zones using practices of defense in depth security. The OT network would treat levels 4–5 as untrusted connections as they were often managed by IT groups and outside connections into OT networks could not be trusted. While the IDMZ was originally conceived to bridge the OT networks (L0–L3) and the IT networks (L4–L5), it did not account for cloud connectivity. As such enterprises struggle to either route data from the lower levels to the IDMZ for cloud connectivity or open an increasing number of ports in the lower levels to facilitate that connectivity while potentially compromising the integrity of the perimeter-based security model.

III. INFLUENCE OF INDUSTRIAL INTERNET OF THINGS AND CLOUD

The push from the Industrial Internet of Things (IIoT) has placed a significant emphasis on the importance of data within an industrial automation and enterprise architecture. As compute surfaces became physically smaller and the cost barrier decreased, intelligent devices within industrial automation became inherently more capable to present and process data. Within increased compute located lower within the industrial automation architecture, additional context, data presentation, data shaping, and data analysis could be performed closer to devices. Additionally, with this additional compute, devices could now expose datasets that were not previously accessible within legacy implementation as the data access and analysis rates were simply not available over the network.

Looking at industrial automation systems from a pure data perspective allows one to provide a construct of "data citizenship" within the system. Data citizenship refers to the ability of components of and industrial automation system to participate in shared data patterns and access mechanisms throughout the system. Traditional data flows within architecture in industrial automation were very regimented and flowed vertically through the ISA-95 model.

A simple example of this is an adjustable speed drive (ASD) and temperature sensor which were required to manage a pumping application within a process. The data flow from a data citizenship perspective would flow from level zero, in this case, the ASD, to a level one PLC. Additional context could be added to data at level one before being elevated to level two and level three systems. An example here may be data to populate an HMI graphic, or long-term persistence data base. This scenario can be seen in Fig. 2.

This traditional method of data collection, aggregation, and persistence has been used under ISA-95 models as standard practice since inception of the standard. The approach though has downsides when constructing an automation controls system. In the scenario shown in Fig. 2, the controller becomes the bottleneck of the system. In other words, all data must flow though the controller, even if the data are not being used to make a control decision. An example of this may be an analog value to be shown on an HMI graphic. Furthermore, controller memory is typically limited and the short-term storage of multiple sources of data may not be the be use of the finite resource.

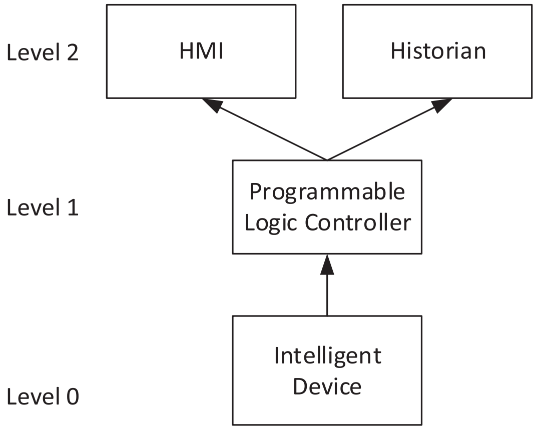

FIGURE 2. Traditional data flow.

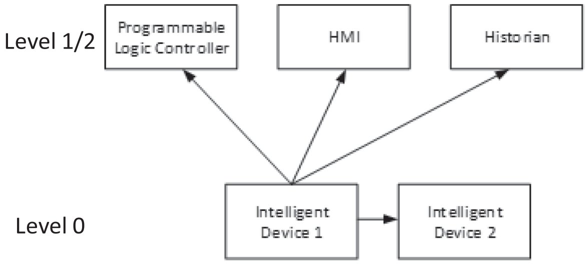

FIGURE 3. Modern data flow.

In terms of "data citizenship," the intelligent device in this scenario has only one client to source data, and that data are then indirectly moved to various layers of the system. This approach limits the line of sight of the intelligent devices within the system, and arguably, diminishes the value of the intelligent device to source information to other subsystems.

The alternative approach to the traditional method of data flow for intelligent devices is seen in Fig. 3. In this method, the intelligent device raised the level of its "data citizenship" due to the ability for the device to self-describe a data model. As the intelligent device can fully describe its data and capabilities, its value increases in the overall architecture from purely a level zero device, providing insights to multiple layers of the architecture, including peer devices.

In Fig. 3, Intelligent Device 1 has expanded its "data citizenship" to four clients: PLC, HMI, Historian, and peer Intelligent Device 2. In this example, the device serves only

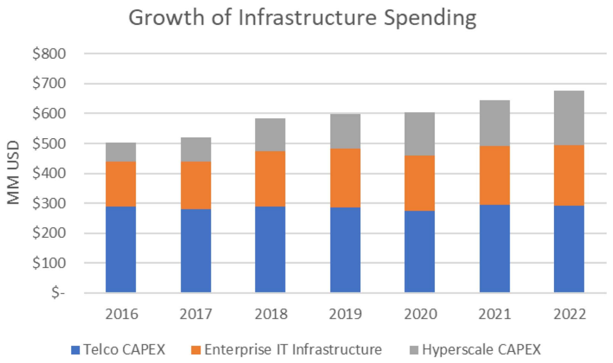

FIGURE 4. Growth of infrastructure spending.

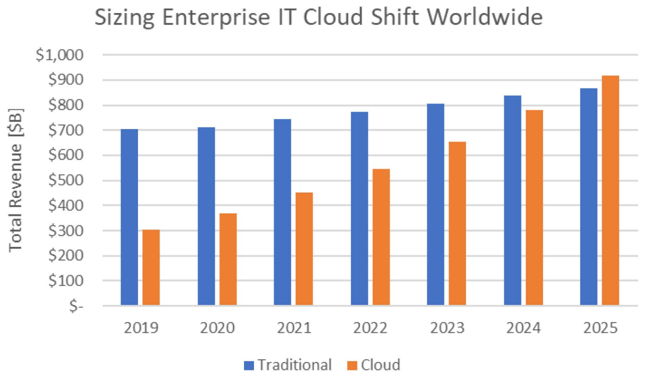

FIGURE 5. Sizing of IT shift to cloud.

data with context to each client that is relevant. For example, the HMI would need graphics data, the historian would require data for long-term trending, and Intelligent Device 2 would inquire about state data from Intelligent Device 1. This example implies the value of the intelligent device to the system and minimizes the need for all information to flow though the singular controller.

For example, an intelligent device may describe data in categories, such as identity, state, runtime, maintenance, and sustainability. The "data citizen" approach leverages the concepts of distributed computing and how it can be leveraged at various locations in an architecture to provide a robust solution that meets the needs for clients. While this approach deviates from traditional ISA-95 models and established best practices, it illustrates the value of data sharing that might cross the established boundaries of the ISA 95 Model. Throughout the document, influences on the enterprise, such as IIoT solutions, and new architecture driven by the hyperscalers will be discussed which will require changes to the way OT networks are configured and managed.

IV. INCREASED INFLUENCE OF IT AND HYPERSCALERS ON CONVERGED IT/OT SYSTEMS

With the changing dynamic of the workforce via the retirement of the "baby boomer" generation and the "mass resignation" event, competent resources with expertise and background to run OT facilities are few and far between. As a result, IT is being thrust in to fill the gap left via retirements and resignations to define infrastructure and standard practices within converged IT/OT solutions.

The influence of the IT organization on OT cloud integration has seen a shift toward the hyperscaler architectures of major players, such as Microsoft, Amazon, Google, and Alibaba. Fig. 4 represents a study conducted across enterprise IT globally over the past 6 years investigating trends in IT expenditures within OT environments [8].

Looking across IT spend over the last 6 years, it can be seen that traditional telecommunications capex has decreased 16 points while hyperscale capex in major corporations have increased 16 points during the same period [9]. Also notable is the stagnation of enterprise IT infrastructure during this same time period. In essence, this means that IT organizations are shifting away from on-premise infrastructure to host environments, and making investments within hyperscaler cloud platforms [10].

Fig. 5 depicts the typical spend by IT between 2019 and forecasted to 2025. This article performed by Gartner was an analysis of major industrial manufactures looking a corporate IT spend and projected IT spend into the future on cloud. It can clearly be seen within Fig. 5 that there has been a significant and steady increase in spend on cloud services by companies surveyed. An inflection point is projected to be reached in 2025 when cloud spend will outpace traditional on premises IT expenditures [10].

The article further discusses the motivating factors for shifting from on-premises to cloud. Major motivating factors for shifting to hyperscaler cloud were noted as elasticity of having resources available when needed, ability to quickly scale, and ability to quickly deploy and provision resources [9].

It is very clear that hyperscalers have had significant influence regarding both technology choices and buying behavior within IT departments. This influence is now being reflected upon OT networks and departments. The authors of this article feel the three most prominent trends being pushed by hyperscalers within OT networks are as follows: hybrid cloud architecture where public cloud, private cloud, and on-premises resources are managed as a single system, distributed computing, typically implemented using Kubernetes, and zero-trust architecture (ZTA). These three items have both significant upsides to revolutionizing how OT networks operate, but directly challenge current OT industry standards and best practices.

V. HYBRID CLOUD

Hybrid cloud architecture is a combination of public cloud, private cloud, and on-premises infrastructure that allows organizations to deploy applications and workloads in the most suitable environment. In industrial automation environments, where data are constantly generated from various sources, such as sensors, machines, and control systems, hybrid cloud architecture will enable organizations to efficiently process, store, and analyze this data in real time.

As cloud (public and private) and on-premises resources become a single distributed system, concepts of edge and cloud begin to flatten out. The hybrid architecture by design is intended to allow workloads to run where they are best suited. Often the article is better suited to run closer to the source of the data, rather than sending it to a central location for processing. This is particularly important in industrial automation environments where real-time processing of data is critical for ensuring safety and reliability. By processing the data locally, organizations can reduce latency, minimize network congestion, and improve overall system performance while still making data available to the broader enterprise for less time critical workloads.

Hybrid cloud also introduces new security models that leverage the security capabilities of both public and private clouds. For example, sensitive data can be stored on-premises or in a private cloud, while less sensitive data can be stored in a public cloud. This approach helps to minimize the risk of data breaches and ensure that critical systems remain secure and available. In addition, by leveraging the security capabilities of cloud providers, organizations can benefit from the latest security technologies and best practices without having to invest in expensive security infrastructure themselves. Over time it is expected that more of those security technologies available from the cloud providers will extend to management of on-premises resources as well, improving overall management, observability, and compliance of systems spanning across both IT and OT.

VI. DISTRIBUTED COMPUTING IN INDUSTRIAL OT NETWORKS

Distributed computing is defined as a form of computing in which data and applications are distributed among disparate computers or systems but are connected and integrated by means of network services and interoperability standards so that they function as a single environment [11]. From a strict definition, distributed computing makes all computers in the cluster work together as if they were one computer. Some of the benefits of this approach are as follows:

1) Scalability: Distributed computing clusters are easy to scale via a scale-out architecture, where higher workloads can be handled by adding new hardware, or scale-up architecture, where higher workloads can be handed by expanding existing hardware.
2) Performance: Distributed computing can leverage parallelism in which each computer in the cluster simultaneously handles a subset of a workload. As a result, the cluster can achieve higher levels of performance.
3) Resilience: Distributed computing clusters typically replicate data across all computer servers to ensure there is no single point of failure. This methodology ensures there is not loss of data due to computation or power failure.
4) Cost-effectiveness: Distributed computing can leverage low-cost hardware, enabling users to provide a variety of solutions, leveraging economies of scale for rollout.

From an industrial automation perspective, the IT/OT convergence, reality places more emphasis on the concepts of where, how, and when to use distributed computing. Major points of emphasis to consider when considering adoption of distributed computing within OT networks are as follows:

- Edge vs. cloud.
- Managed vs. unmanaged.
- Real time vs. nonreal time.
- Always connected vs. sometimes connected.
- On-premises vs. cloud.
- Secure vs. “Partially” secure.
- Governed vs. nongoverned.

All the points of emphasis for the consideration of distributed computing within modern converged IT/OT architectures are important and impact the overall performance and security of the system. The next section of this article will discuss the use of Kubernetes to manage these distributed resources.

VII. KUBERNETES

Kubernetes is a container orchestration platform that automates deployment, scaling, and management of containerized applications. First established by Google in 2015, Kubernetes morphed into an opensource project owned and maintained by the Cloud Native Computing Foundation (CNCF). The original version of Kubernetes was designed with an intent of managing large cloud-native workloads which saw large adoption within data centers. With the project’s transfer to CNCF, Kubernetes received backing by major hyperscalers as a deployment and management option for containerized workloads [12].

The cloud community became interested in Kubernetes due to its capability to provide built in commands for deploying applications, rolling out changes at scale to applications, scaling applications elastically to support changing needs and requirements, and monitoring applications. Perhaps, more importantly, it was recognized as a platform that could be used with or without cloud providers to create highly effective distributed compute systems. Research has shown that the Kubernetes system is capable of meeting the needs of cloud applications for both performance and management needs [12], [13], [14], [15], [16].

Kubernetes is an extremely capable solution for container orchestration. It is valued in the community of practice both for its high availability and its elasticity of scale to dynamically needs loading needs. One downside to this flexibility is the amount of infrastructure that is needed to support a full-blown Kubernetes architecture.

While the centralized management, high availability and scalability of Kubernetes is generally desirable, it is not always practical. Smaller offerings, such as K3s and Micro K8s, have been developed to enable solutions of small device

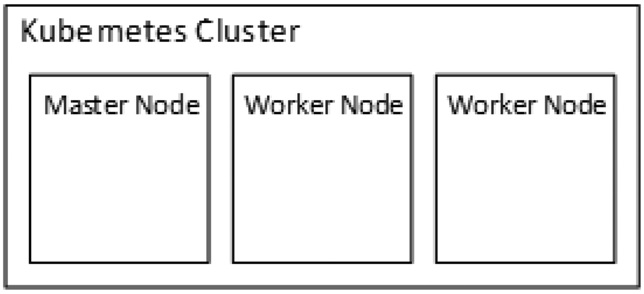

FIGURE 6. Kubernetes architecture.

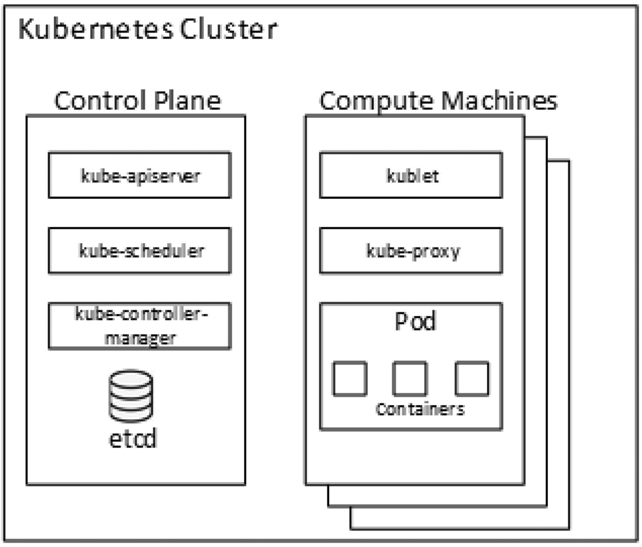

FIGURE 7. Node architecture.

clusters, and single node clusters that often target IoT and IIoT scenarios [17].

These “small” implementations of Kubernetes are designed to run on small and sometimes standalone devices, while still providing the exact same tooling for configuration, management, and observability simplifying the management of distributed compute systems made up of a variety of resources [16].

A. KUBERNETES ARCHITECTURE

Kubernetes deployments leverage a collection of nodes. These nodes can be servers that are physical or virtual, on premise or cloud native. In the Kubernetes architecture, there are two major types of nodes: master node and worker node as depicted by Fig. 6.

Worker nodes are where applications execute within a pod. Fig. 7 defines the basic architecture of a worker node. Each worker node contains three important components: the container runtime, a kublet, and a kube-proxy. The container runtime is responsible for the execution of the container. The kublet is the agent responsible for communicating with the Kubernetes control plane on the master node. This component is responsible for executing commands received from the master node as well as a hypervisor for the running pods. The kube-proxy is the network proxy on each node and serves as the network broker of communications inside and outside the cluster [14].

The master node runs the Kubernetes control plane. In order to provide high availability, multiple master nodes can be run within a cluster [16]. There are four main components within the master node as shown in Fig. 7: kube-apisever, kube-scheduler, kube-controller-manager, and the etcd database. The kube-apiserver acts as the frontend communications to the control plane [15]. The kube-scheduler is responsible for the scheduling of pods on worker nodes. It must consider resourcing requirements as well as user-imposed requirements. The kube-controller-manager acts as a supervisor to the current state of the system matches the desired configuration settings. When using a managed Kubernetes offering from a cloud provider, the kube-controller-manager provides an interface between the cluster and a cloud provider. Finally, the etcd database is a distributed key-value store for all cluster related data is housed. This database is only accessible via the API server [13].

Kubernetes is an extremely capable solution for container orchestration. It is valued in the community of practice both for its high availability and its elasticity of scale to dynamically needs loading needs. One downside to this flexibility is the amount of infrastructure that is needed to support a full-blown Kubernetes architecture.

While the centralized management, high-availability and scalability of Kubernetes is generally desirable, it is not always practical. Smaller offerings, such as K3s and Micro K8s, have been developed to solutions of small device clusters, and single node clusters that often target IoT and IIoT scenarios [17].

These “small” implementations of Kubernetes are designed to run on small and sometimes standalone devices, while still providing the exact same tooling for configuration, management, and observability simplifying the management of distributed compute systems made up of a variety of resources [16].

VIII. KUBERNETES MANAGED RESOURCES

While Kubernetes is made up of nodes participating in the cluster, there are cases when there are resources which are unable to host the kublet and be an active participant in the Kubernetes cluster. New patterns and technologies are emerging that enable resources outside of the cluster to be managed (indirectly) by the cluster. One emerging technology is an open-source project called Akri [18].

Akri is a service plugin introduced into the Kubernetes cluster containing two primary components: the Akri Controller and the Akri Agent. The Agent is effectively responsible for discovering resources (or the absence of known resources), once discovered the Controller handles the creation or deletion of Pods and Services that enable interaction between resources in the cluster and the discovered devices.

The Pods and Services created by the Akri Controller are intended to act as “Proxy Services” to the discovered resources [19]. These proxy services can be configured to perform any number of actions, such as data access, configuration interfaces, as well as participation in a ZTA by introducing policy enforcement points (PEPs) to devices that might not otherwise be able to participate in ZTA.

IX. ZERO TRUST WITHIN KUBERNETES

While Kubernetes itself does not create a ZTA, it does provide a platform capable of implementing various principles discussed in Section X as follows:

1) Identity and access management (IAM) can be enforced by using Kubernetes through integrations with available identity providers, such as active directory, or OAuth2 providers.
2) Principle of least privilege can be enforced by using Kubernetes role-based access control) to limit access to resources based on the user’s role.
3) Operating on the assumption all networks and devices are untrusted can be enforced by using network policies to control communication between the deployed pods within the cluster [20].
4) Monitoring and logging can be provided through various integrations, such as Kubernetes Audit Logs or Prometheus (an open-source monitoring system), which allow you to monitor all the activities within the cluster.
5) Kubernetes enables the creation of a layered defense approach by allowing for multiple layers of security to be implemented. The ability to scale the cluster by adding new components allows integration with third-party tools for intrusion detection or vulnerability scanning. Kubernetes also allows for network segmentation to enable different security levels for different workloads.

The following section of this article will discuss the importance of trust in modern IT and OT networks and how modern security principles are impacting the design and considerations for OT networks interfacing to cloud infrastructure.

X. TRUST IN MODERN IT/OT NETWORKS

As the OT infrastructure becomes more interconnected with the balance of the enterprise, it is natural that IT will take a more significant role in the management of OT infrastructure. To reduce the risk of cybersecurity attacks and control the impact of a breach that was not avoided, new approaches to security are needed. The perimeter-based models associated with ISA-95 will no longer be enough. The interconnectivity driven by data demands, remote workers, connected partners, and customers have changed the model from a series of hierarchical networks owned/managed by a single entity to a graph of systems and networks without a single point of management [21].

NIST provides the following operative definition of ZT and ZTA:

ZT provides a collection of concepts and ideas designed to minimize uncertainty in enforcing accurate, least privilege per-request access decisions in information systems, and services in the face of a network viewed as compromised. ZTA is an enterprise’s cybersecurity plan that utilizes ZT concepts and encompasses component relationships, workflow planning, and access policies. Therefore, a zero-trust enterprise is the network infrastructure (physical and virtual) and operational policies that are in place for an enterprise as a product of a ZTA plan [22].

ZT and ZTA take a different approach than found in ISA-95. When the ISA-95 model was conceived, contemporary concepts, such as least privileged necessary were not commonplace, let alone evaluated at the per-request level as advocated by ZTA today. Rather resources implicitly trusted each other, generally only limited by the network perimeter that contained them. This trust often did not require any specific authentication, let alone specific authorization to determine the level of access to grant a given request.

Where ISA-95 assumes everything inside a given network perimeter is secure, ZT assumes the network is already compromised. Forcing the implementor to not only consider how to control access to resources, but how to contain lateral movement once a resource or network has been compromised. The use of security zones provides the implementor additional logical boundaries requiring both authentication and authorization before a request can traverse to resources in the targeted zone. ISA-95 by contrast may limit traffic between zones, however, if a node that has access to traverse zones becomes compromised, it is likely to breach zones as well.

Supporting the operational definition of ZT and ZTA above NIST has developed the following tenants:

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 dynamic policy—including the observable state of client identity, application/service, and the requesting asset—and may include other behavioral and environmental attributes.
5) The enterprise monitors and measures the integrity and security posture of all owned and associated assets.
6) All resource authentication and authorization are dynamic and strictly enforced before access is allowed.
7) The enterprise collects as much information as possible about the current state of assets, network infrastructure, and communications, and uses it to improve its security posture. [9]

While NIST provides high-level tenants, such as those above to define a ZTA, they do not go so far as to prescribe specific technology or implementations. In fact, it is often possible for adopters to enhance their existing policies and

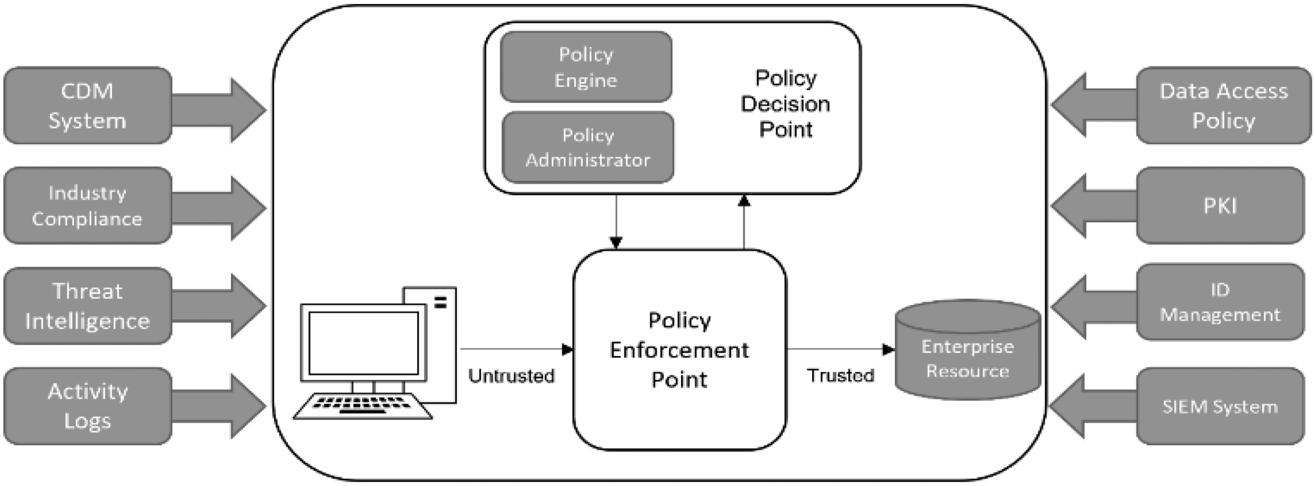

FIGURE 8. Zero-trust architecture.

practices to achieve the goals of ZT with little or no additional technology expenditures.

XI. COMPONENTS OF A ZERO-TRUST NETWORK

A ZTA involves many components. While many of these components and the roles they play will be new to the OT network, most modern IT systems are already using some or all the components suggested in Fig. 6 below.

The components within Fig. 6 are as follows.

A. POLICY ENGINE

This component grants or denies access to a resource for a given request. While it uses enterprise policy for this decision, best practices suggest taking additional input from external sources, such as threat intelligence services and databases as in input to a trust algorithm intended to create an overall score to not only determine if access should be granted or denied, but in extreme circumstances it may be determined that all access to either the requesting or targeted resource should be revoked.

While the policy engine (PE) makes and logs decisions about access to a given resource, the enforcement of that decision is the job of the policy administrator (PA).

B. POLICY ADMINISTRATOR

Based on the decision made by the PE, the PA creates or destroys session-specific communication channels between the subject (requestor) and the target resource. This is done by communicating to the relevant PEPs. In a sense the PA can be thought of as a “broker’ that is responsible for passing the appropriate authentication and authorization to a targeted PEP for the specific session. If the session is denied or if the approval is revoked, the PA instructs the PEP to close the connection. In some implementations, the PE and the PA may be combined.

C. POLICY ENFORCEMENT POINT (PEP)

This component can be thought of as a proxy or a gatekeeper for communication to a given resource. When commands from the PA are received, the PEP will either create or destroy a connection to the targeted resource. Throughout the session, the PEP is also expected to monitor the traffic to the targeted resource and close the connection with the intended request is completed.

XII. DATA INPUTS FOR ZERO-TRUST POLICY

As pictured in the Fig. 8, ZTA requires several inputs to be successful. Policy rules as well as external and environment data are combined to allow the PEs to make access control decisions with the most complete information possible. The following sections provide a high-level summary of the inputs to a ZTA architecture as defined by NIST.

A. CONTINUOUS DIAGNOSTICS AND MITIGATION (CDM)

The continuous diagnostics and mitigation (CDM) system maintains an inventory of the digital assets in the organization including the current state of the asset regarding software and configuration versions. The CDM is often the system responsible for deploying updates to those same assets. In many cases, the CDM also maintains an inventory of enterprise-approved and/or nonapproved components and if an asset has any known vulnerabilities.

This information becomes a direct input to the PE allowing the PE to make more informed decisions including information about the software and configuration of assets making or responding to requests.

B. INDUSTRY COMPLIANCE

Industry compliance refers to the regulatory policies that the enterprise must adhere to. Industry-specific regulations for healthcare, finance, as well as general information security and privacy are all examples requiring specific policy to ensure adherence to the regulation. The industry compliance system may be a standalone system, or in many cases will simply be the framework for policy generally used but specifically including required regulatory guidelines.

C. THREATS MODELING AND LOGGING

The threat intelligence feed is intended to inform the PE of information from multiple sources (internal and external) that provide information about newly discovered vulnerabilities. External data may include newly discovered flaws in software and new malware. Internal data may include log information aggregated from various systems including information about reported attacks to other assets or systems that the PE will want to isolate from the rest of the enterprise.

D. DATA ACCESS POLICIES

Data access policies are the foundation for authorizing access to a resource. They provide the definition of basic access rights for accounts and applications in the enterprise. These base policies should be built on the core roles of the organization and then augmented with more applications specific and dynamic data generated by the PE.

E. ENCRYPTION AND IDENTIFICATION MANAGEMENT

Typically built on x.509 certificates, the enterprise public key infrastructure (PKI) is the system used to generate and log certificates issued by the enterprise to resources throughout the enterprise. This can include applications, devices, and various other subjects and services. These certificates are often integrated with the global certificate authority ecosystem.

F. ID MANAGEMENT

The ID management system is responsible for storing and managing the user accounts and identity records for users in the enterprise. This system stores appropriate data for all users (name, email address, certificates, etc.) used to uniquely identify each individual user for authentication. Additional information regarding role(s) and access attributes, associated or assigned assets may also be included in the ID management system. The ID management system may utilize the PKI system for artifacts such as x.509 certificates associated with the user account.

G. SECURITY INFORMATION AND EVENT MANAGEMENT

Security information and event management (SIEM) system are the point of aggregation where security centric information from various systems is collected and made available for later analysis. This data can be integrated with the PE and/or used to refine and improve existing policies. In some cases, data from the SIEM can be used to report to external sources and warn of possible attacks.

Unlike traditional approaches to network security, ZT incorporates significantly more granular and comprehensive tooling. Inclusion of external data, e.g., CDM, industry compliance system, threat intelligence feed (s), network and system activity logs, and data access policies, to enhance the PE is incorporated to provide a more robust view of the system and possible threats. This has the potential to dynamically change an access decision based on awareness of a vulnerability discovered in the device or platform requesting access.

The NIST specifications listed above are by designed to be technology and provider agnostic. Various vendors have created specific approaches to ZT and ZTA.

XIII. IMPLEMENTATION RISKS AND CHALLENGES

In today's connected world, network security poses a considerable challenge due to the increasing prevalence of cybersecurity threats and attacks. Organizations are adopting zero-trust security models to protect their networks. This transition, however, has risks associated with it that organizations must be aware of, and take steps to mitigate as they make their transition.

A. ADDRESSING OUTDATED OPERATING SYSTEMS AND FIRMWARE

Outdated operating systems and firmware can contain vulnerabilities that can be exploited by attackers to gain unauthorized access to systems and data. This could result in data breaches, theft of sensitive information, and compromise of critical systems and networks. Organizations must keep their operating systems and firmware up to date to mitigate these risks.

B. BALANCING COST CUTTING AND NETWORK SECURITY

Years of cost cutting can result in a lack of investment in security solutions, which could result in making the network more vulnerable to attacks. This can cause a lack of resources to detect and respond to security incidents, and a shortage of trained personnel to manage security operations. Organizations must balance cost cutting measures with the need for adequate security measures to protect their networks.

C. MITIGATING SECURITY RISKS OF WEAK, DEFAULT, INSECURE, OR NONEXISTENT PASSWORDS

Using no passwords, default passwords, or insecure passwords can allow attackers to easily gain unauthorized access to systems and data. This can lead to data breaches, theft of sensitive information, and compromise of critical systems. Organizations must implement strong password policies and enforce the use of complex and unique passwords to mitigate these risks, across all accessible assets.

D. ENSURING A COMPREHENSIVE ASSET INVENTORY

Lack of an inventory of assets can make it difficult to detect and respond to security threats, as organizations will not know what they need to protect. This could result in a lack of visibility into the security posture of the network and a lack of control over sensitive data. Organizations must maintain an up-to-date inventory of all assets, including hardware, software, and data, to effectively manage their network security, and consider the scope of policies in relation to this asset inventory.

E. MANAGING COMPLEXITY AND RESOURCE CONSTRAINTS

Implementing a ZTA can be complex and resource-intensive, requiring integration of multiple technologies and solutions, a clear understanding of technology requirements, a budget that allocates adequate resources, and technical expertise to integrate components, such as IAM, multifactor authentication, and SIEM. Organizations must also have the resources to maintain the solution, including investments in hardware, software, and personnel training. These challenges can be particularly difficult for organizations with limited resources.

F. MANAGING FALSE POSITIVES

ZT solutions can generate a high volume of alerts, many of which may be false positives. Organizations must have the resources in place to manage these alerts effectively, and to ensure that genuine security threats are not missed. Also, they must have the technical expertise to differentiate between false positives and genuine security threats and must have a response plan in place to address genuine threats. False positives can also lead to a decrease in user trust and adoption of the zero-trust solution, in addition to the costs and technical challenges of managing them.

G. INTEGRATING ZERO-TRUST SOLUTIONS IN EXISTING SYSTEMS

Integrating a zero-trust solution with an organization's existing systems and networks can pose a significant challenge, especially in brownfield environments where existing systems are often outdated or use proprietary technologies that may not be easy to update. In such environments, organizations must have the technical expertise to integrate the zero-trust solution with their existing systems, which could be a time-consuming and costly process.

In greenfield environments, where an organization is starting from scratch, implementing a zero-trust solution is easier as there is no need to integrate with existing systems. However, organizations must still ensure that the zero-trust solution is compatible with their future technology infrastructure and can scale as the organization grows.

In both greenfield and brownfield environments, organizations must carefully assess the impact of implementing a zero-trust solution on their existing systems and networks and allocate adequate resources and personnel to ensure a successful implementation.

H. EVOLVING IT POLICIES AND MANAGING UNKNOWN RISKS

Implementing a zero-trust solution represents a significant shift in network security, and organizations must be prepared to maintain this new security paradigm over time. This includes regularly reviewing and updating their security policies, regularly testing, and updating their security solutions, and investing in ongoing training for security personnel to ensure they have the necessary expertise to effectively manage network security. Failure to maintain a strong focus on network security can result in vulnerabilities that attackers can exploit, leading to security breaches and data loss. It is important for organizations to have a plan in place for continuously monitoring and updating their network security to address emerging threats and take advantage of new technologies.

XIV. CASE STUDIES

The following section defines multiple case studies of various types of manufacturing that have implemented cloud technologies to assist their operations. All references to individual companies have been removed to eliminate commercialism.

A. FOOD MANUFACTURER

A major food manufacture uses data from controllers, drives, and other sensors to monitor the state of equipment and throughput of lines. The architecture of this solution involves a combination of software from automation vendors, hyperscaler cloud providers, and custom reporting solutions created by the companies IT department.

Challenges have included connectivity to a variety of different automation systems, inconsistent implementations of security, and data interfaces as well as the need to normalize data. The outcome has been increased visibility into production, allowing for better coordinated maintenance schedules, improved uptime, and improved ability to meet retail customer's just in time requirements.

B. MEDICAL EQUIPMENT MANUFACTURER

A manufacturer of medical equipment has developed robotic fabrication machines, which use a combination of PLCs, variable speed drives, and robotics to fabricate various prosthetics and other medical implementations. IoT technology is employed to allow for the remote monitoring of the equipment (by the manufacturer) to ensure the machines service level agreement is being met. If necessary automated work orders may be generated to dispatch a service team to correct on site problems.

Challenges encountered with this solution are primarily around the disparity of the environment the machines are deployed in. Unlike a typical factory floor setting, these machines may not be subject to traditional Purdue network models but must still adhere to the (often more modern) security standards of the hospital or medical organization employing the machine.

C. AUTOMOTIVE MANUFACTURER

As with many large multisite enterprises, the automotive manufacturer has a combination of decentralized maintenance (resources at each plant) with a centralized group of higher skilled senior resources in headquarter type locations to support the remote sites. After introducing IoT technology to surface information about automation device firmware versions, runtime data, and predictive maintenance information, the centralized team was able to reduce travel time to remote locations and the company saw an over reduction in unplanned downtime.

D. WAREHOUSING LOGISTICS COMPANY

A warehousing and logistics company utilized cloud technology to fleet manage their assets throughout their enterprise. As a hybrid system consisting of multiple manufacturers equipment, creates data visibility and collection issues.

One challenge faced is the collection of enterprise-wide data to allow data science easy access to provide asset performance with standard organizational benchmark. Much of this data had to come directly from devices as it was not easily available within other levels of the control system. This approach flattened the ISA 95 models as defined previously in this article.

A second major challenge this organization faced was the ability to fleet manage configurations for devices, inclusive of firmware and security patches within the system. This challenge was due to equipment manufactures adjusting configurations during service and not setting the device configurations back to the corporate standard. As a result, corporate engineering would need to be dispatched cost time and extra money to reset these configuration errors. Utilizing cloud technology, all assets can be centrally managed and configurations backed up, monitored, and pushed to maintain the fleet of equipment.

XV. SUMMARY OF APPROACHES DESCRIBED

This section summarizes the pros and cons of both the ISA-95 Model vs. ZT approaches.

A. PROS PERIMETER-BASED SECURITY (PURDUE NETWORK MODEL)

1) Clear network boundaries: Perimeter-based security models establish well-defined network boundaries, making it easier to manage and monitor security within those boundaries.
2) Segmentation of IT and OT systems: The Purdue model helps segregate IT and OT systems, reducing the risk of exposure and improving overall security.
3) Familiarity and historical success: The Purdue model has been used for many years in industrial control systems and is familiar to many security professionals.

B. CONS PERIMETER-BASED SECURITY (PURDUE NETWORK MODEL)

1) Assumes trust inside the perimeter: Once an attacker breaches the perimeter, they can move laterally within the network, potentially causing significant damage.
2) Less effective with cloud and edge computing: The increasing adoption of cloud-based services and edge computing has made it difficult to maintain a well-defined perimeter.
3) Difficult to adapt to modern threats: The Purdue model struggles to cope with dynamic and sophisticated security threats that are increasingly common in today's interconnected world.

C. PROS OF ZERO-TRUST SECURITY MODEL

1) No inherent trust: The zero-trust model assumes that no user, device, or service is to be trusted by default, reducing the risk of lateral movement within the network for an attacker.
2) Granular control: ZT allows organizations to apply security policies at the workload level, giving organizations more control over their applications, data, and interactions.
3) Better suited for modern IT architectures: ZT is designed for cloud, edge computing, and remote work environments, which are becoming increasingly prevalent in modern times.

D. CONS OF ZERO-TRUST SECURITY MODEL

1) Complexity: Implementing a zero-trust model can be complex, requiring a deep understanding of the organization's systems, resources, personnel, and security requirements to design and manage it effectively.
2) Resource-intensive: It may be challenging for organizations with limited resources (time, budget, and personnel) to successfully implement and maintain a zero-trust security model.
3) False positives: A high number of alerts may be generated by the zero-trust solution, requiring resources and expertise to manage and differentiate genuine threats from false positives.

XVI. CONCLUSION

Having served Industrial Automation for several years, the ISA-95 model may be on its way to obsolescence. New technology, such as cloud and IIoT, driven by an ever-increasing demand for data have introduced new risks into the once secure perimeter-based network system. As enterprise systems have expanded to the cloud and now back to the edge creating hybrid distributed systems, the landscape for IT/OT integration continues to evolve. New approaches to architecture that spans from IT into OT is being influenced by nontraditional players. Rather than automation vendors or network providers, now it is the hyperscaler cloud provider, partnered with corporate IT who is driving the change to the OT landscape.

The multilayer hierarchy of networks in ISA-95 is beginning to flatten out. IDMZs are being replaced with software-defined networks and policy-based access control. Rather than controlling access to a network, the new goal is to control access from workload to workload which might be ephemeral, running in a virtualized environment on a cluster of physical computers.

Technologies, such as Kubernetes and containers, are becoming ubiquitous, running an increasing number of workloads reliably in both the cloud and on premises, but the OT systems they are trying to integrate with were built with no understanding of these technologies. In many cases, those OT systems have little or no means to perform basic IAM themselves and as a result require a higher level abstraction to perform these duties if they are to be integrated into a modern network architecture. Using technology, such as ZT PEPs, to control access to automation devices for instance, may help bridge the gap between yesterday's automation technology and the requirements and expectations of today's network architecture but that is only one small part of an overall approach to replacing ISA-95.

Moving forward, users of automation technology will be relying on their automation vendors and cloud providers to work together to create validated reference architectures that not only allow for new technology to solve for these new requirements, but more importantly reduce the risk of migrating brownfield, existing implementations to new architectures to meet the demands of the next 10 years or more.

XVII. FUTURE RESEARCH

The authors plan to continue research in cloud technologies as applied to industrial automation. It is believed that the market is now at an inflection point where cloud will have a larger influence than ever on industrial manufacturing.

Immediate next steps for research will be to devise a design of experiment with more ZT principles and provide additional insights and information in this space. Additional research will be conducted on AKRI and how automation devices can be self-discovered and drive actions within a Kubernetes environment.

REFERENCES

[1] Y. Bobbert and J. Scheerder, "Zero trust validation: From practice to theory: An empirical research project to improve zero trust implementations," in *Proc. IEEE 29th Annu. Softw. Technol. Conf.*, 2022, pp. 93–104, doi: 10.1109/STC55697.2022.00021.

[2] Y. Ge and Q. Zhu, "Trust threshold policy for explainable and adaptive zero-trust defense in enterprise networks," in *Proc. IEEE Conf. Commun. Netw. Secur.*, 2022, pp. 359–364, doi: 10.1109/CNS56114.2022.9947263.

[3] ANSI/ISA-95.00.01-2010 (IEC 62264-1 Mod) Enterprise-Control System Integration, I. S. o. America, 2010. [Online]. Available: https://www.isa.org/products/ansi-isa-95-00-01-2010-iec-62264-1-mod-enterprise

[4] G. Rathwell, "ISA-95—Setting the stage for integration of MES and ICD systems," in *Proc. Inst. Eng. Technol. Seminar Enterprise Integr. Control Syst.*, 2006, pp. 113–114.

[5] B. Wally, C. Huemer, and A. Mazak, "Entwining plant engineering data and ERP information: Vertical integration with AutomationML and ISA-95," in *Proc. 3rd Int. Conf. Control, Automat. Robot.*, 2017, pp. 356–364, doi: 10.1109/ICCAR.2017.7942718.

[6] D. C. Mazur, R. A. Entzminger, P. A. Morell, J. A. Kay, and E. Syme, "Defining the industrial demilitarized zone and its benefits for mining applications," *IEEE Trans. Ind. Appl.*, vol. 52, no. 3, pp. 2731–2736, May/Jun. 2016, doi: 10.1109/TIA.2016.2530045.

[7] C. Gifford, *When Worlds Collide in Manufacturing Operations: ISA-95 Best Practices Book 2.0*. Pittsburgh, PA, USA: ISA, 2011.

[8] "2022 capex analysis—Growth in hyperscale and enterprise spending; telco remains in the doldrums," S. R. Group, 2023. Accessed: Jan. 30, 2023. [Online]. Available: https://www.srgresearch.com/articles/2022-capex-analysis-growth-in-hyperscale-and-enterprise-spending-telco-remains-in-the-doldrums

[9] A. Weissberger, "Synergy research: Growth in hyperscale and enterprise IT infrastructure spending; telcos remain in the doldrums," 2022. [Online]. Available: https://techblog.comsoc.org/2023/01/27/synergy-growth-in-hyperscale-and-enterprise-it-infrastructure-spending-telcos-remain-in-the-doldrums/

[10] M. Haranas, "Cloud provider spend on IT capex climbs as telecom falls," 2023. [Online]. Available: https://www.crn.com/news/cloud/cloud-provider-spend-on-it-capex-climbs-as-telecom-falls

[11] M. Raynal, "A look at basics of distributed computing," in *Proc. IEEE 36th Int. Conf. Distrib. Comput. Syst.*, 2016, pp. 1–11, doi: 10.1109/ICDCS.2016.109.

[12] N. T. Nguyen and Y. Kim, "A design of resource allocation structure for multi-tenant services in Kubernetes cluster," in *Proc. 27th Asia Pacific Conf. Commun.*, 2022, pp. 651–654, doi: 10.1109/APCC55198.2022.9943782.

[13] A. P. Ferreira and R. Sinnott, "A performance evaluation of containers running on managed Kubernetes services," in *Proc. IEEE Int. Conf. Cloud Comput. Technol. Sci.*, 2019, pp. 199–208, doi: 10.1109/CloudCom.2019.00038.

[14] S. Dan, "Managing Kubernetes clusters," in *Official Google Cloud Certified Associate Cloud Engineer Study Guide*. Hoboken, NJ, USA: Wiley, 2019, pp. 175–207.

[15] M. M. Khalel, M. A. Pugazhendhi, and G. R. Raj, "Enhanced load balancing in Kubernetes cluster by minikube," in *Proc. Int. Conf. Smart Technol. Syst. Next Gener. Comput.*, 2022, pp. 1–5, doi: 10.1109/ICSTSN53084.2022.9761317.

[16] S. Telenyk, O. Sopov, E. Zharikov, and G. Nowakowski, "A comparison of Kubernetes and Kubernetes-compatible platforms," in *Proc. 11th IEEE Int. Conf. Intell. Data Acquisition Adv. Comput. Syst.: Technol. Appl.*, 2021, pp. 313–317, doi: 10.1109/IDAACS53288.2021.9660392.

[17] M. Fogli et al., "Performance evaluation of Kubernetes distributions (K8s, K3s, KubeEdge) in an adaptive and federated cloud infrastructure for disadvantaged tactical networks," in *Proc. Int. Conf. Mil. Commun. Inf. Syst.*, 2021, pp. 1–7, doi: 10.1109/ICMCIS52405.2021.9486396.

[18] *Project-AKRI*. CNCF, 2023. Accessed: Mar. 31, 2023. [Online]. Available: https://github.com/project-akri/akri

[19] R. Vaño, I. Lacalle, P. Sowiński, R. S-Julián, and C. E. Palau, "Cloud-native workload orchestration at the edge: A deployment review and future directions," *Sensors*, vol. 23, no. 4, 2023, Art. no. 2215, doi: 10.3390/s23042215.

[20] D. D. Silva and D. D. Ambawade, "Building a zero trust architecture using Kubernetes," in *Proc. 6th Int. Conf. Convergence Technol.*, 2021, pp. 1–8, doi: 10.1109/I2CT51068.2021.9418203.

[21] A. Wylde, "Zero trust: Never trust, always verify," in *Proc. Int. Conf. Cyber Situational Awareness, Data Anal. Assessment*, 2021, pp. 1–4, doi: 10.1109/CyberSA52016.2021.9478244.

[22] "Zero trust architecture," N. I. o. S. a. Technology, 2020, doi: 10.6028/NIST.SP.800-207.

CYRIL PERDUCAT (Member, IEEE) received the master's degree in engineering and the second master's degree in international project managementin from the Ecole Nationale Supérieure des Arts et Métiers, Paris, France, and the ESCP Europe Business School, Paris, France, respectively, in 1994.

He is currently a Senior Vice President and a Chief Technology Officer. He and his team of technologists are focused on growing Rockwell's already broad technology offering and expanding what is possible in the industrial automation space. He is responsible for charting the company's technology roadmap, our reusable intellectual property library, and research related to core technology. Throughout his career, he has focused on developing innovative solutions to enhance manufacturing technologies. He has an extensive background as a technology strategist, particularly in systems, software, networks, and solutions businesses. He joined the Rockwell following a 25-year career with the Schneider Electric, where he most recently served as an Executive Vice President in Internet of Things and Digital Offers. He has a passion for innovating and building teams, and a drive for implementing successful outcomes based on effective change management.

DAVID C. MAZUR (Senior Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering from the Virginia Polytechnic Institute and State University, Blacksburg, VA, USA, in 2011 and 2012, respectively, and the Ph.D. degree in mining engineering for his work with automation and control of the IEC 61850 standard from the Virginia Polytechnic Institute and State University, in 2013.

He is currently working as a Senior Manager with the Rockwell Automation, Milwaukee, WI, USA, with a current focus on digital experiences for industrial automation products. His experience includes application development in heavy industry automation and infrastructure.

Dr. Mazur is an active member of the IEEE IAS and serves as working group chair for the Communication-Based Protection of Industrial Applications Working Group. He also serves as a member of the Mining Industry Committee as well as the Industrial and Commercial Power Systems Committee. He is also an active voting member of the IEEE Standards Association.

WES MUKAI (Member, IEEE) received the B.Sc. degree in industrial engineering from the Stanford University, Stanford, CA, USA, in 1989, and the MBA degree in harvard general managment from the Harvard Business School, Boston, MA, USA, in 1995.

He is currently working as a Chief Engineering Officer/VP with the Rockwell Automation, Milwaukee, WI, USA, with a focus on user experience, software architecture and engineering across data, edge, analytics, cloud, and industry solutions. He brings a wealth of experience in the high-tech industry across the networking, computing hardware, telecommunications, and software industries for companies, such as Teradyne, PwC, Infogear, Cisco Systems, SAP, and General Electric Transportation.

SCOTT N. SANDLER (Senior Member, IEEE) studied business information systems and accounting with the University of Phoenix, Phoenix, AZ, USA, 1997–2002. He has worked in Industrial Automation Software space for more than 25 years, with a focus on Information Solutions and Cloud Computing. During this time, he has held many positions related to design, development, architecture, and product management of various products and solutions.

MICHAEL J. ANTHONY (Member, IEEE) received the B.Sc. degree in computer and electrical engineering from the Marquette University, Milwaukee, WI, USA, in 2008, and the master's degree in business administration from the Milwaukee School of Engineering, Milwaukee, WI, USA, in 2021. He is currently working toward the Ph.D. degree in manufacturing systems focused on communication technologies with the Capitol Technology University, Laurel, MD, USA.

He started with the Rockwell Automation as a Software Development Engineer, in 2005, on a variety of information and HMI focused products in the FactoryTalk portfolio. He has held roles as a Product Manager for a variety of HMI, communication, and security software products in the Rockwell Automation portfolio. He is currently focused on applications communication technology in the Rockwell Automation Strategic Development organization in the office of the CTO.

JON A. MILLS (Senior Member, IEEE) received the B.Sc. degree in computer science from the Ohio University, Athens, OH, USA, in 2012.

He started with the Rockwell Automation, in 2013, with a focus on integrating intelligent devices into industrial control systems. He, working as a Software System Architect, continues to focus on device integration within both traditional OT networks, but also with respect to the OT/IT boundary.