Workgroup: Independent Submission
Internet-Draft: draft-km-iotops-iiot-frwk-00
Published: 16 January 2022
Intended Status: Informational
Expires: 20 July 2022
Authors: K. Makhijani L. Dong
Futurewei Futurewei

Framework For Integrated Industrial Networks

Abstract

Industry control networks host a diverse set of non-internet protocols supporting Industrial-IoT and legacy device connections. These networks are physically separated from the enterprise networks and have been slow to adopt the virtualization technologies. Virtualization is necessary to remove the boundaries between the enterprise and process control networks. This document specifies a framework for the converged industrial network. Specifically, it focuses on the virtual PLC scenario. It covers transition technologies required for the convergence of industrial devices with the enterprise application endpoints.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 20 July 2022.

Table of Contents

1. Introduction

There is a little cross-over between the network technologies used in the Operational Technology (OT) and Information Technology (IT) environments as the architecture of industrial networks has evolved independently from the IT networks. While industrial networks focussed on deterministic communication, safety, and reliability of process control, they have trailed behind in the adoption of virtualization capabilities. Virtualization is necessary for the industry automation to enable compute and data-intensive services at scale. Moreover, it is also a necessary tool for the convergence of OT and IT systems.

Although the virtualization of SCADA and other systems (HMI, MES, Historian, etc.) has happened, the low-level PLCs have remained on the factory floor. Virtualization of PLC is a topic of great interest for fully automated and remote operations in the manufacturing and similar process control industry because it allows direct control of low-level processes from the applications.

This memo studies the overall network requirements for support of virtualized PLCs and identifies their characteristics. The benefits include:

- *Flexibility to control the devices from application-level logic thus improving automation,

- *Potentially eliminate need for dedicated PLCs on the floor that reduces interconnection and integration overheads.

- *Ability to leverage high-end general purpose processing platforms to perform complex compute intensive operations.

- *Adapt rapidly to changing business requirements with software, avoiding hardware changes.

Enabling PLC virtualization imposes a set of challenging requirements. Broadly they can be viewed from

1. PLC perspective: the mechanisms with which it integrates with other business applications while preserving PLC logic, and its real-time or deterministic constraints when communicating with the devices it interacts with;

2. Network perspective: the impact on the network when PLCs are not directly connected to devices i.e., requirements to reliably move data between the virtualized PLC elements and OT devices (sensors, and actuators) while maintaining operational safely.

This document presents the baseline industrial architecture in Section 3 and demonstrates that the current hierarchical architecture poses difficulty for the adoption of PLC virtualization in industrial networks Section 4. Section 5 further develops approaches to the support for virtualized PLCs.

A distributed conceptual model that could potentially address the above limitations is presented for discussion as a converged industrial-network architecture Section 6. Finally, a summary of requirements is extracted in Section 7.

This document discusses those requirements and proposes a path to converged industrial networking.

2. Terminology

Industrial Control Network:
The industrial control networks are interconnection of equipments used for the operation, control or monitoring of machines in the industry environment. It involves different level of communications - between field bus devices, digital controllers and software applications

Industry Automation:
Mechanisms that enable machine to machine communication by use of technologies that enable automatic control and operation of industrial devices and processes leading to minimizing human intervention.

Control Loop:
Todo

Feedback Control Loop:
Todo

Programmable logic controllers (PLC):
Industrial computers/servers for the control of manufacturing processes such as assembly lines.

Supervisory Control and Data Acquisition (SCADA):
Software System to control industrial processes and collect and manage data.

Distributed Control Systems (DCS):
Systems of sensors and controllers that are distributed throughout a plant.

Manufacturing Execution System (MES):
Systems that connect production equipment across the factory floor, or multiple plants or sites.

Fieldbus Devices:

Operational Technology field devices include valves, transmitters, switches and actuators etc.

Integrated Industrial Network (IIN):

The term introduced in this document to represent a converged view of OT and IT networks. Virtualized PLC (vPLC):

A software component of PLC, in which the control part of factory devices is decoupled from the I/O component. With vPLCs, the I/O stays local to the machines (sensors, actuators, and drives), while the controller logic lives as a software service implemented over RT- hypervisors.

2.1. Acronyms

*HMI: Human Machine Interface

*MES: Manufacturing Execution System

*CIN: Converged Industrial Network

*IIC: Industrial Internet Consortium

*IDMZ: Industrial Demilitarized Zone

*PLC: Programmable Logic Controller

*PDU: Protocol Data unit

*SCADA: Supervisory Control And Data Acquisition

*DCS: Distributed Control System

*OT: Operational Technology

*IT: Information Technology

3. Industrial Network Architecture

The physical network architecture for process control and operations centric networks as shown in Figure 1 is rigidly hierarchical. Note that the figure is over-simplified and in general, each level will have additional hierarchies to extend networks for scale. For example, a PLC that is controlling a group of fieldbus devices may be controlled by another PLC controller which runs ProfiNet protocol. In such cases protocol translation gateways are needed. Between these gateways there may exist a set of intermediate network switches to extend the range (physical distance) and scale (number of devices) of connectivity on the factory floor. Similarly, system

integrators also need a variety of translation gateways to extract and integrate data from field devices as an input to MIS, HMI and other enterprise applications.

The hierarchical architecture comprises of security oriented zones that are combined together to represent ICA-95 model (or Purdue model see Appendix A) in which each zone further contains well-defined levels. The communication across the zone tends to get complex as each zone runs over a different technology. Among the three zones (Manufacturing, IDMZ and Enterprise), the enterprise zone network is all IP while manufacturing and IDMZ network on factory floor are a combination of IP and Industrial protocols. IP-based routers are used in manufacturing-zone when there is need to extend the network across different cells on factory floors. There are a large number of IP based firewalls and gateways that perform translation in IDMZ but are required to look at transport payload to determine the industrial protocols and corresponding matching rules.

Higher layer applications directly interact with PLCs for send and receive commands and data from the field devices. The data generated by sensors is transmitted by PLCs to industry control systems (SCADA, HMI, MES). Both DCS and SCADA systems collect data from process instruments and respond with commands to the actuators. They control several process control loop instances simultaneously to handle complex processes. Originally, such systems were built using proprietary hardware and software, operating in its own zone without additional connectivity. Operators had to work from centralized control room because these systems did not support remote access. The best practices for data delivery to other systems was in the form of reports, which caused significant time lag.

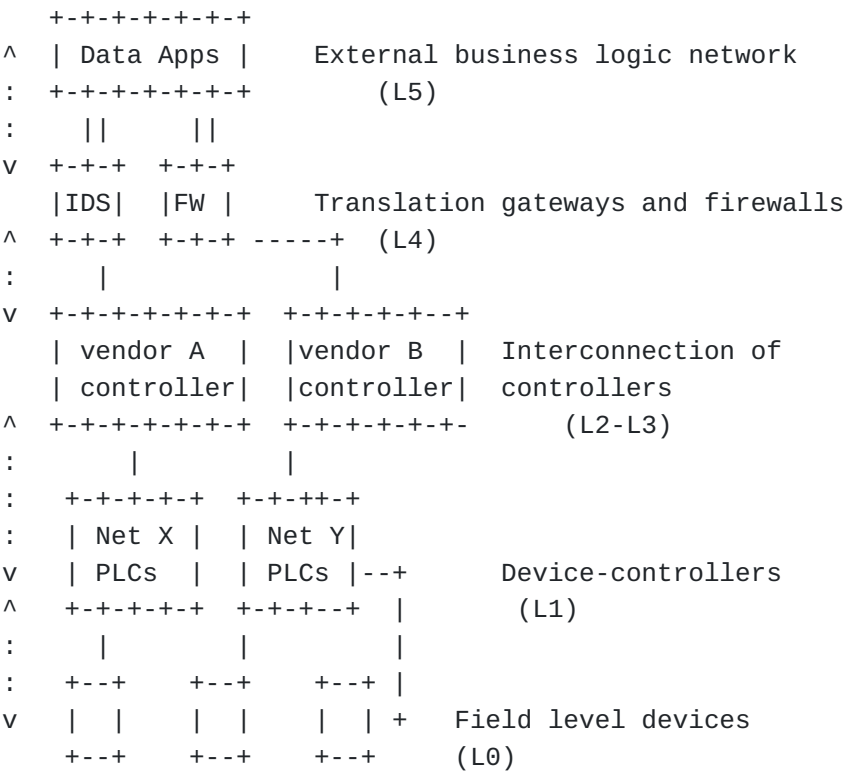

Figure 1: Hierarchy of Functions Industrial Control Networks

I/O devices (sensors, actuators) generate a large volume of data and also accept process control commands. For an application to handle a variety of low-level protocol translations can be extremely challenging, therefore, solutions such as OPC-UA [OPC_ARCH] or common messaging broker mechanisms MQTT [MQTT_SPEC] are deployed. While OPC-UA is a common representation of data collected from different I/O devices; MQTT is messaging service designed for systems with low resources. Both are higher-level protocols that are generally transmitted over TCP as shown in Figure 2 below.

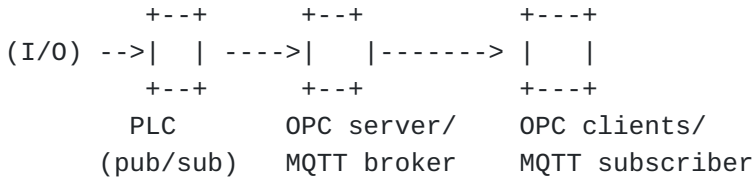

Figure 2: Protocol agnostic data collection in Industry Networks

4. Challenges and Limitations

As is evident from the ICA-95 model (described in Appendix A), the business applications are centralized in the enterprise networks. In this network architecture, with PLCs virtualized, they may be colocated at the edge of manufacturing zone, with the supervisory

control systems, or in Enterprise zone along with business applications.

Alternatively, virtualized PLCs enable new capabilities such as utilizing cloud to edge-aware network architectures by flexible placement of applications and allocation of resources to different components in industry control systems. Virtualzed PLCs serving from the edges can meet latency constraints to close the control loop for process automation.

4.1. Associating virtualized PLCs with I/O Devices

A physical PLC is generally associated with a few I/O devices and is directly connected. The I/O modules are not required to authenticate or perform any verification on connection. As virtualized PLC may be on the other side of the network, the I/O device requires an authentication mechanism to connect to the PLC. This is necessary to maintain reliability and safety of the system and prevent unauthenticated PLC to interact with the software.

4.2. Expectations from network performance

Virtualiaton allows consolidation of compute, storage, and network resources, as well as independence from custom hardware. The magnitude by which compute capability is improved allows a single virtualized controller to handle more complex and faster scan cycles.

> Note: A scan cycle is generally the time taken to read the inputs, execute the program (e.g. ladder logic), and update the outputs. The actual scan time is affected by the processing speed of the PLC, the size of the program, the type of instructions used in the program. Therefore, in virtualized PLCs general-purpose processor speed and the amount of memory available is much higher than most physical PLCs.

Then, the performance of the network to handle communicaton delays, packet formation, processing and forwarding overheads become critical to overall system performance. Additionally, use of edge-compute platforms is expected for both consolidating resource consumption and lowering the operational costs.

4.3. Multiprotocol supporting PLCs

Another difference between physical and virtualized PLC is that with virtualization of PLC, a single controller can communicate with a different group of I/O devices over one or more non-internet protocols such as Modbus, Profibus, CANbus, Profinet [SURV], etc. Each of the protocols specifies its packet format.

The benefit is a reduction in the number of controllers, but the requirement challenge is to provide a standard communication format for different I/O devices. Since it is not feasible to communicate packets in native (Fieldbus protocol) format due to address scale limitations (field bus devices have limited address space up to 256 devices), a network element or an end-device is required to perform some format translation.

As an example, a factory floor is composed of different cell sites. A set of PLCs controls I/O modules or machines in each cell. Traditionally, when there is an inter-connection requirement between two or more cells, the protocol translation is carried out between the cells. With physical PLC, the translation would be done at the controller, but with virtualized, it may require translation capability on the network element connecting to I/O devices or the devices themselves. Moreover, additional deployments are not integrated with the existing network, creating a new network.

4.4. Identification of virtualized PLC

The fieldbus devices are serial buses and identify PLC as a device with a specific bus address. In the large scale network and in the application layer this much information is insufficient. It may be required for virtualized PLC to support dual addresses, one exposed for the I/O module and other for IT applications.

Moreover, as an exmaple, it is no longer sufficient to indicate basic address 0x14; it may require to specify 'device 0x14, of cell - C1 and factory floor, F1, PLC bus address 0x1 in communication path. The reachability to a specific I/O module should have complete information from virtualized PLC.

4.5. Security Aspects

The fundamental paradigm of security as expressed in ICA-95 architecture changes with virtualized PLC since the PLCs are now moved away from the local manufacturing zone. The zone based security design considerations can employ either or both of the approaches:

1. Describe mechanisms to abstract the manufacturing and other zones. This maybe achieved using secure communication channel approaches such as such as VPN, IPSEC etc. In this approach traffic admission policies applied on a PLC, will now be applied on traffic entering virtual PLC.

2. Describe mechanisms for location-specifc (site) perimeter security. This maybe achieved using conventional firewall methods.

When zone-based secuity rules are leveraged, location-specific security policies may be more coarse grained. For example, an ingress firewall rule will be required to verify and authenticate that the source site is permitted to send traffic for specified destination.

5. Evolving Networks for virtualized PLCs

This section conceptualizes a fully virtualized industrial control system in which all the components - network, compute, storage, and applications run on virtual platforms to better understand the gaps and requirements.

5.1. Virtualization of Components in Industrial Systems

Virtualization enables the separation of software from hardware. It is the foundation for disaggregating different system components such as operating systems, management tools, service logic, and data.

In the case of industrial control systems, by virtualizing SCADA, MES, and HMI, etc. [VPLC_CONV], these softwarized systems are run on commodity hardware or general-purpose CPUs. Main benefit of virtualizing supervisory and control systems is that the overall cost can be controlled since specialized hardware is not required, while operators can perform software upgrades more frequently. Such virtualized software can be placed anywhere, often close to the source of data it needs to process. This, in turn, leads to leveraging edge compute networking for multi-site integration.

While applications and services are beginning to get disaggregated, PLCs' virtualization is very early stage. Conceptually, a virtual PLC means that the controller functions are separated from the I/O modules of the devices.

5.2. Incremental Approaches

Similar to SCADA, MES, HMIs, virtualized PLC may be located anywhere in industry control architecture. However, expanding beyond a factory cite, requires special security considerations discuss in Section 4.5. Adding new virtualization capabilities may require and overall redesign of the network infrastructure which may not desirable in all the cases specifically, from the perspective of maintaining same level of security, reliability and safety requirements.

Therefore, we envision that the following different approaches are possible:

5.2.1. Softwarized

This is the basic approach with minimal change and minimal impact. A PLC software is virtualized and runs on a commodity hardware supporting legacy interfaces to I/O modules. This type of change is isolated to a specific PLC functionality and the only benefit is use of commodity hardware. Potentially, there is a one to one replacement of physical to software PLC.

5.2.2. Localized

In this approach PLC is virtualized and can run on a commodity hardware as above; additionally providing a clear separation between hardware and software components by relying on protocol translation gateways as part of the network edges that connect to I/O modules conceptually represented in Figure 3.

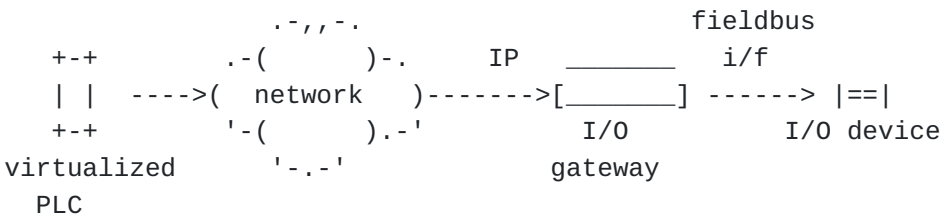

Figure 3: virtualization of PLC and separation from I/O devices

Depending on the compute capabilities of the hardware, different instances of virtualized PLC may run simultaneously or a group of PLCs are bundled together in a single instance of virtualized PLC. Furthermore, a virtualized PLC maybe hosted on the same hardware along with SCADA or ICS components.

There are two new concepts that will need to be formalized: - The I/O translation gateways are a new component in ICA architecture. These are interface translators on an edge network element devices that perform conversion of network side PDU to device side PDU. - identification of the virtualized PLC as discussed in Section 4.4.

The incremental benefit beyond the use of commodity hardware is the ability of encapsulating complex logic in a single instance. A clean separation between PLC logic from I/O module allows changes to PLC logic and I/O devices independently. Since the location of virtualized PLC is with in manufacturing zone there is no impact on the security design.

5.2.3. Distributed

This is the eventual goal to support virtualized PLC in a location independent manner. All benefits considered in Section 5.2.2 apply

with an advantage of leveraging third-party edge-compute infrastructure as a tenant.

However, security zones are impacted as discussed in Section 4.5.

6. Converged Architectural Concepts

Since a virtualized PLC now looks like an IT-centric software component with OT-specific capabilities, the industrial network framework should evolve accordingly to handle virtual entities in the network. This is referred to as a converged network architecture and its conceptual model with significant functions, components and interfaces are discussed in this section.

The foundational concepts of converged industry network architecture has three design principles (i) Ability to virtualize end-points, (ii) Disaggregation of I/O Devices, and (iii) Converged industry-network fabric to support communication between virtualized endpoints and I/O modules.

6.1. Overview

Figure 4 represents a converged network fabric (bottom-left) that enables the transfer of data between software system components (top-left) and the physical devices (bottom-right). The fabric is a shared network infrastructure that allows all the characteristics required in the industrial control networks, such as deterministic, low-latency, and real-time communications.

In Figure 4, "B. factory site" represents one or more cell (locations) of the physical devices. Each cell group belongs to one physical site, and there can be multiple such sites. A cell site may be the smallest network in this fabric.

"A. Software components" emulate different OT and IT functions hosted in a cloud-like environment which is distributed, i.e., components may be located at various sites or at the edge to support low-latency, deterministic applications. Both field and software components are connected over a converged network (shown as "C. <network"> in Figure 4), which presents a unified view of the network infrastructure interconnecting software and field components.

The converged fabric is composed of 3 types of network elements with specific roles.

CISN (Converged Industry Service Node):
The application or service nodes that get virtualized and softwarized maybe instantiated anywhere in the Industry network independent of the I/O module placement. They are placed in cloud

or at the edge or the factory floor itself depending on the usage and type of application. From the communication perspective, these nodes following the state of the art in IT could use technologies and protocol such as IPv6 addresses [RFC8200], and service chains [RFC8300] for steering between different service nodes. Their interface to network will have specific transport requirements (when not transmitted as overlay).

CIFR (Converged Industry Fabric Router):
Network nodes that form the converged fabric and understand the traffic flows between I/O modules and applications. CIFRs are all expected to run a uniform suite of protocols. In a much simplified view the fabric maybe a core IPv4 or IPv6 based forwarding plane, regardless whether overlay technologies (such as VPN, VxLAN etc.) are used). Potentially interconnected over different physical media technologies (commodity or TSN) Ethernet.

CIFRs also perform functions for WAN interfaces and multi site interconnections. The routers performing this function will be at the upper edge of the physical network.

CIIG (Converged I/O-Gateways) These gateways are the lower-level edges of the converged fabric.They are very similar to the existing PLC gateways that are purpose-built for translation between fieldbus/fieldbus or fieldbus/IT protocols. In contrast to traditional gateways CIIGs could be stateless and optionally provide more secure control to I/O device.

6.2. Component Virtualization

In the existing deployments, components such as HMI, Historian, MES, and SCADA systems run on dedicated hardware. Virtualizing these system components can be consolidated on a single general-purpose hardware platform, reducing the number of hardware devices and improving the security of data exchanges among these systems.

6.2.1. Virtualized PLC

A virtualized PLC decouples controller logic from the I/O component. It allows integration of supervisory and control software components as part of the execution environment by leveraging mature IP-based technologies.

Although an exploratory work [VPLC_CONV] and [VPLC_IIC] propose I/O field-buses to be replaced with the high-speed, deterministic media (such as Ethernet). The legacy systems (such as serial fieldbus interfaces) will continue to exist in the foreseeable future. Thus, the architecture must support the communication between the field-bus I/O and PLCs, even when the PLCs are virtualized. This implies

that in some cases the fabric will still need protocol translation gateways on cell sites, but they need to be close to the the I/O modules i.e. at the edge of the converged fabric.

6.2.2. Application and Services

Component virtualization enables co-location of different service functions on the same hypervisor or entirely at a different location regardless of what security zone they belong to. The constraints aware path chain can be set using [RFC7665]. Moreover, it provides multiple service function chain to support different applications. This type of architecture along with NFV [ETSI_GS_NFV_003] can be extremely resource efficient.

Several sensors emit time-series data, that can add to the bandwidth consumption to the information going to the cloud. Deploying big-data application closer on the edge and scale them on-demand provides a sophisticated tool to disaggregate processing of sensory data and summarize for the cloud-enterprise applications.

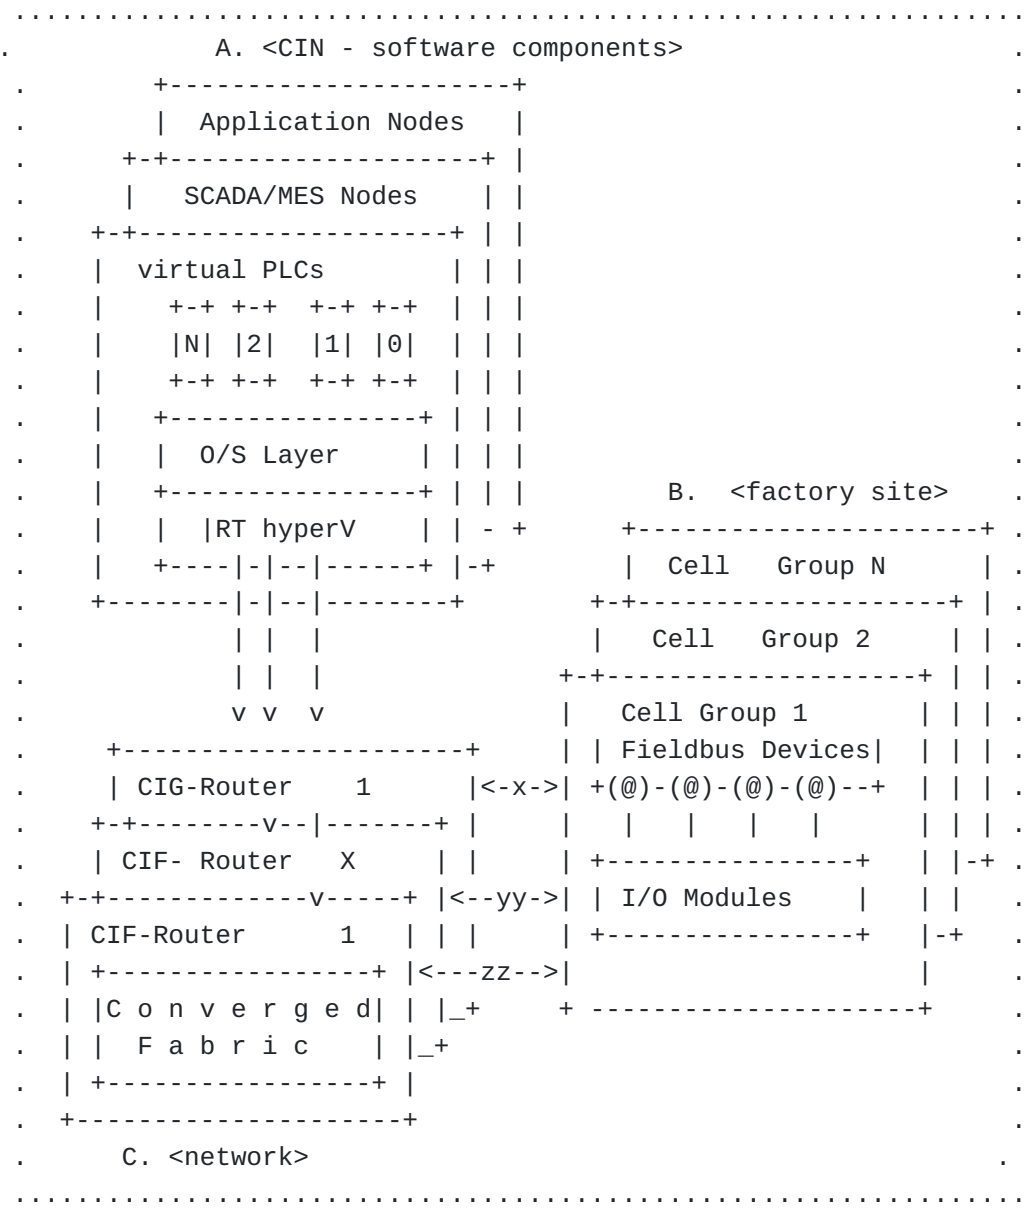

Figure 4: Converged Industrial Network Architecture

6.3. I/O Modules and Field Components

A manufacturing facility can be located at more than one site and each site is further divided into cells. Further the machinery, actuators, sensors are associated with the cell connected to the PLCs. These controllers run different protocols such as Ethernet, RT-Ethernet, Modbus, ProfiNet etc.

In a vPLC supported environment, the I/O cards are responsible for media access conversion from in and out of the converged fabric. Even the support for legacy PLCs is similar to vPLCs, with the role reduced to only translation function.

6.4. Converged Fabric

Converged fabric shown as "B." in Figure 4 is central to the architecture. The connectivity is largely Ethernet based (except I/O device interfaces), potentially running IP protocols on the switches and routers in the network.

Since this is a logical fabric, the connectivity is local on a factory floor and can be extended to multiple sites. Interconnecting different sites will use WAN functions. The fabric breaks the hierarchical structure and topology can now be designed as fat-tree (or leaf-spine) network which provides overall more number and multiple deterministic paths between two end points.

A key characteristic of legacy Industry networks is that they do not require frequent changes and therefore, topology changes are not dynamic. The fabric could potentially use a combination of software-defined connectivity with IP routing protocols. The routing protocols will maintain the infrastructure reachability among the network nodes and software-defined solutions will manage flow of traffic in a deterministic manner addressing the low-latency and deterministic data delivery of certain type of flows.

7. Requirements

7.1. General Requirements

*Basic requirement for converged fabric is the efficiency of connections between the IT software and floor I/O modules or modules, i.e. the connection to a low-level factory devices is uniform (or homogenized) manner. Uniformity implies a variety of endpoints interconnect in an identical fashion without requiring device specific translations. Efficient connections lead to less processing or states in the network with improved resiliency and performance. There maybe opportunities to design packet formats with minimal overheads by using in-band programmability paradigms that carry embedded metadata and control information relating to reachability, latency, jitter, reliability, and exceptions characteristics. This type of approach is expected to reduce configurations and number of policies required for data steering through the network. Existing methods that maybe used, evaluated or extended include IP with TSN, DETNET[DETNET], reachability headers SCHC, IPv6 compression schemes or may be evaluated against newer schemes.

*The converged-fabric shall support traffic segmentation. As connections change between the devices, it should not have adverse effect on deterministic, low latency behavior on the other segmented traffic. Each segmented traffic may be associated

with a different protocol or traffic profile including legacy traffic format and profiles. In case the fabric supports legacy traffic flow or devices, its performance should be no worse than before the fabric. The methods to support segmentation include virtual network technologies inside the fabric such as VxLAN, VPNs, etc.

*The fabric protocols used must not limit to a constrained physical topology. It should support efficient multi-path distributed connectivity framework to prevent bottlenecks, traffic concentration. Even in the industrial networks growth in data-generation is obvious as more number of telemetry and reporting sensors are being added in the system. Managing bandwidth for different types of data (operational, control, statistics) should be considered. It is expected that an entire cell may be added or removed on-demand. This type of changes shall be dynamic (i.e. does not require long-term planning) and non disruptive (no impact on existing performance metrics) to other traffic-paths in the network. These characteristics scale better with layer-3 network designs.

*The fabric is a logical representation of LAN and WAN connections. The traffic ICA-95 zone-transfer may happen anywhere in the logical topology. In this regard, appropriate perimeter or zone admission policies MUST be designed and enforced in such a manner that are not bound to a specific location. Alternately, it may require two-stage policies first one to validate traffic conforms to the zone policies and second conforming to specific service behavior. A potential approach here could be the use of semantic addressing [I-D.draft-farrel-irtf-introduction-to-semantic-routing] where part of the addresses may describe match rules for zones and services.

7.2. Device Specific Requirements

Device functions and operation does not change. The requirements here are related to how that are reached, identified and discovered in the network.

*Addresses scope: As the scale of industry network grows, there will be many same type of devices with limited address space (a fieldbus or ModBus address limits up to 256) all across the floor. Therefore, a structured addressing scheme is necessary to uniquely identify each device from the operator's command center.

*Converged Namespace: Each industry vertical could have different preferences for how it chooses to view devices and applications in the system. It should be possible to identify all the endpoints as part of a system defined namespace. The solution

should not require different operations and management schemes for industry I/O modules vs IT applications. A common namespace that aligns with business goals can simplify management. For example, assigning segmented identifiers for each level (PLCs, cell sites, type of application etc.) and concatenating them together provides helps industry operations considering that factory devices do not change their location often in the topology.

*Network Identifiers: Each device should be identifiable in the network by what application it can talk to. The network identification should be provided for setting up security or firewall policies. Note: today legacy devices do not have network identifiers. With virtual instances of PLCs, it is to be determined how different instances of the same PLC will be identified, discovered and associated. Moreover, it maybe desirable to support variable length identifiers to handle both IT servers and I/O module type devices.

*Legacy support: Architecture must provide legacy device support for deployed protocol formats and their core capabilities. This is needed to maintain non-disruptive operations.

*Once part of the larger fabric, the devices must be discoverable. Given the 'physical-location' of the device can often be preprogrammed, device on-boarding should allow a device to be auto-configurable.

*The auto configuration procedures should be efficient, i.e., comparable to the processing capabilities of the I/O devices.

*On-boarding procedures (manual or automatic) must have built-in or well-defined authentication procedures.

*Device policies: Each device connects to at least one controller (PLC or a gateway). When the network detects a misbehaving controller, the policies should define the default behavior of the device (such as quarantined from the network along with the gateway or allow it to operate autonomously with default settings, or shutdown etc).

Further motivation and analysis for adapting to OT/IT asymmetric address formats is covered in [I-D.draft-km-industrial-internet-requirements].

7.3. Key Performance Indicator Requirements

*Performance of industry operations depend on the deterministic behavior of devices. The network must preserve and support this attribute such as the legacy device connections with controllers.

*Safety mechanisms : To keep a factory floor hazard and accident free environment, the system should implement built-in mechanisms for proper operation of a devices (i.e. software commands sent from vPLC must not exceed thresholds).

*Security: mechanisms should be implemented to protect man-in-the-middle attack. Knowing that E2E encryption procedures on device can impact low-latency, due to low processing power, light-weight mechanisms should be devised (such as reduce the data to be encrypted, account it in KPIs, etc).

7.4. Virtualization Related Requirements

The topologies in the manufacturing zones do not change frequently and devices are designated in a zone or a cell for long-term use. Such observations can help simplify network designs. As such, industry networks substantially benefit from a hybrid approach of software-defined networking and distributed routing. Former for initial provisioning, latter for reachability and health of the fabric. Such hybrid approaches eliminates the need for complex routing protocol features.

*Edge compute and networking - TBD.

7.5. Virtualized PLC Requirements

*Solution must provide a secure method of pairing, authenticating a virtualized PLCs with their I/O devices.

*Virtualization allows multiple PLCs to control the same device. This can potentially lead to conflicts in device operation. Therefore, careful policies are required to prioritize operation across the PLCs.

*Currently, L0 and L1 devices (ICA_95) do not use any transport protocol. The data is embedded after control header. With a network layer solution, TCP maybe too heavy for field-bus devices. Some other means of assuring device delivery will be needed.

8. IANA Considerations

This document requires no actions from IANA.

9. Security Considerations

The architecture at the very least must adhere to the security guidance provided by ICS-95.

10. Acknowledgements

11. References

11.1. Normative References

[RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <https://www.rfc-editor.org/rfc/rfc7665>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/rfc/rfc8200>.

[RFC8300] Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed., "Network Service Header (NSH)", RFC 8300, DOI 10.17487/RFC8300, January 2018, <https://www.rfc-editor.org/rfc/rfc8300>.

11.2. Informative References

[DETNET] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <https://www.rfc-editor.org/rfc/rfc8655>.

[ETSI_GS_NFV_003] "Network Functions Virtualization (NFV) Architectural Framework", 2017.

[I-D.draft-farrel-irtf-introduction-to-semantic-routing] Farrel, A. and D. King, "An Introduction to Semantic Routing", Work in Progress, Internet-Draft, draft-farrel-irtf-introduction-to-semantic-routing-02, 14 January 2022, <https://datatracker.ietf.org/doc/html/draft-farrel-irtf-introduction-to-semantic-routing-02>.

[I-D.draft-km-industrial-internet-requirements] Makhijani, K. and L. Dong, "Requirements and Scenarios for Industry Internet Addressing", Work in Progress, Internet-Draft, draft-km-industrial-internet-requirements-00, 10 June 2021, <https://datatracker.ietf.org/doc/html/draft-km-industrial-internet-requirements-00>.

[IIC] "Industry IoT Consortium", n.d., <https://www.iiconsortium.org>.

[IIC_TALK] William Diab, W., "Overview of IIC – Building the IIoT Ecosystem", 12 October 2021, <https://github.com/iot-dir/

Meetings/blob/main/20211012/slides/Diab_IIC_Overview_for_IETF_1021_rev2.pdf>.

[ISA95] "ANSI/ISA-95.00.01-2010 (IEC 62264-1 Mod) Enterprise-Control System Integration - Part 1: Models and Terminology", n.d., <https://www.isa.org/standards-and-publications/isa-standards/isa-standards-committees/isa95>.

[MQTT_SPEC] "MQTT Version 3.1.1 Plus Errata 01", December 2015, <http://docs.oasis-open.org/mqtt/mqtt/v3.1.1/mqtt-v3.1.1.html>.

[OPC] "Open Platform Communications", n.d., <https://opcfoundation.org>.

[OPC_ARCH] "OPC 10000-1 - Part 1: Overview and Concepts", 2 November 2017, <https://opcfoundation.org/developer-tools/specifications-unified-architecture/part-1-overview-and-concepts/>.

[OPC_INFO] "OPC-UA Information Model Specifications", n.d., <https://opcfoundation.org/developer-tools/specifications-opc-ua-information-models>.

[SURV] Galloway, B. and G. Hancke, "Introduction to Industrial Control Networks", IEEE Communications Surveys & Tutorials Vol. 15, pp. 860-880, DOI 10.1109/surv.2012.071812.00124, 2013, <https://doi.org/10.1109/surv.2012.071812.00124>.

[VPLC_CONV] Cruz, T., Simoes, P., and E. Monteiro, "Virtualizing Programmable Logic Controllers: Toward a Convergent Approach", IEEE Embedded Systems Letters Vol. 8, pp. 69-72, DOI 10.1109/les.2016.2608418, December 2016, <https://doi.org/10.1109/les.2016.2608418>.

[VPLC_IIC] Lou, D., Graf, U., and M. Tseng, "Virtualized Programmable Logic Controllers. An Industrial Internet Consortium Tech Brief", 7 September 2021, <https://www.iiconsortium.org/pdf/IIC-Edge-vPLC-Tech-Brief-20210907.pdf>.

Appendix A. Appendix A. Purdue Model (ICA-95)

The International Society of Automation (ICA) has developed a model [ISA95] to describe automated interfaces between enterprise and control systems. In this widely deployed hierarchical model, five levels are defined and they follow a strict ordering of interfaces across the levels. At the lowest level 0, are the physical devices

while enterprise applications are at level 5. In between these two levels, there are several supervisory, management, and intermediate data collection applications that provide information to

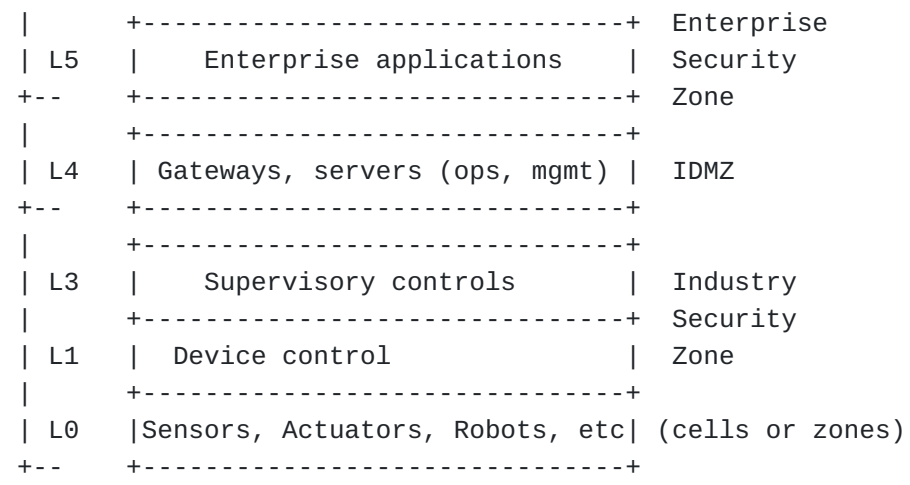

Figure 5: ISA 95 or Purdue model of Automation Pyramid

A.1. Separation between Manufacturing and Enterprise Networks

The ICA-95 architecture recommends hierarchy, thereby a separation between factory devices and applications through three different security zones called Manufacturing, DMZ and enterprise zones as shown in Figure 5 as below:

* **Enterprise Security Zone:** The IT applications reside in enterprise networks and perform tasks necessary for business operations such as inventory control, supply-chain logistics, schedule and capacity planning. They need to collect data from the OT systems in order to make those decisions.

* **Industrial Demilitarized Zone:** The OT and IT networks were designed to prevent direct communication between them. The IDMZ serves as an information sharing layer between the IT and OT (L4 and L3) systems. This indicates that additional security rules, inspection and protection of device identity and access is necessary when transiting from L3 to L4.

* **Manufacturing Zone:** Consists of Levels 0 through 3 site wide production system. Operations at level 3 (L3) Support site-wide view of the production system. They also provide data to L4. Area supervisory control (L2) performs operation and control over a zone or smaller area in a production floor. Each area has specific set of tasks or operations to perform. Basic control at level 1 (L1) is for the actual control of the equipment. The L1 components such include PLCs; they send commands to L0 equipments to perform tasks (e.g. start motor,

alter pressure level, or reduce motor speed). Finally, actual process takes place at level 0 (L0). At this level for the process equipments performing actual operations are performed. This include equipment and devices such as motors, pressure valves, temperature, speed, etc sensors, etc.

The devices or controllers at level 1 are the ones of specific interest for virtualization and the corresponding challenges are covered in later section.

A.2. Collaborating with SDOs with Industry Network Focus

The paradigms of networking in OT are quite different than IP based best-effort networking protocols. Yet, IETF protocols are extensively used in OT applications. Often, it is not possible to get contributors directly from the OT sectors, then it would make more sense to coordinate with well-established consortia where OT scenarios and requirements are is discussed may be utilized. Two well established foundations are IIC [IIC] and OPC-UA [OPC]. For example, a [IIC_TALK] provided overview of IIC activities.

Industrial IoT Consortium (IIC) provides use cases, scenarios, and best-practice frameworks to solve specific problems and solution pain points. It is a rich resources of case studies and demonstrations of different test beds. The IIC itself is not involved in standards development, but may help in formalizing requirements, further insights into solutions developed in IETF, and potentially help adoption of those solutions.

Open Platform Communications-Unified Architecture (OPC-UA) provides interoperability across different hardware platforms using a standard data model. It standardizes various information models, corresponding client-server architecture and defines necessary access mechanisms to those information models. The OPC-UA is an abstraction layer to provide common interface to different data look-up and event notifications. A number of information models are provided by OPC-UA can be found here [OPC_INFO]. Foe example, OPC has a specification on PLCs. It abstracts PLC specific protocols (such as Modbus, Profibus, etc.) into a standardized interface allowing HMI/SCADA systems to interface with a middleware that converts generic-OPC read/write requests into device-specific requests and vice-versa.

Note: OPC-UA information model similar to YANG?

IETF solutions will focus on leveraging or extending IETF technologies for IT and OT integration which is at the infrastructure or communication layer. Thus, providing protocols that could potentially benefit higher-level OPC-UA work.

Both IIC and OPC could provide guidance to the lower level work.

*For Discussion: assuming there is an IIN framework - how does it fit in the OPC-UA architecture and facilitate adoption of existing information models.

Authors' Addresses

Kiran Makhijani
Futurewei
Santa Clara, CA 95050,
United States of America

Email: kiran.ietf@gmail.com

Lijun Dong
Futurewei
Santa Clara, CA 95050,
United States of America

Email: lijun.dong@futurewei.com