

# Virtualized Programmable Logic Controllers

## A Paradigm Shift Toward Industrial Edge and Cloud Computing

An Industrial Internet Consortium Tech Brief

20210907

**Authors:**

**David Lou, Ph.D.**

Chief Researcher, Huawei Technologies

*zhe.lou@huawei.com*

**Ulrich Graf**

Senior Engineer, Huawei Technologies

*ulrich.graf@huawei.com*

**Mitch Tseng, Ph.D.**

Managing Member, TSENG InfoServ, LLC

*mitch@T-infoserv.com*



---

Table of Contents

**1 The Integrated, Edge-Based Architecture..... 3**  
**2 Implementation Considerations ..... 4**  
**3 A Look into the Future..... 7**  
**4 Business Considerations ..... 7**  
**5 Related Presentations in IIC..... 8**  
**6 About the Authors ..... 8**

Table of Figures

**Figure 1-1. The manufacturing industry is adopting Edge Computing for IT/OT convergence...4**  
**Figure 2-1. Automation platform architecture .....7**

---

Most industrial automation applications demand high reliability and availability for control devices and associated input and output (I/O) elements. The first discrete Programmable Logic Controller by Modicon (now Schneider Electric) was introduced in 1968 and Allen-Bradley, in 1971, coined the term PLC. Since then PLCs have been widely adopted as the means of control in production lines in the manufacturing industry. Although they generally employ an array of PLCs to execute I/O controls precisely, each PLC needs communication ports and a controller unit, making it bulky and expensive. It is also expensive to update programs once deployed.

Virtualized PLCs (vPLCs) on commercial-off-the-shelf (COTS) infrastructure can replace a large number of individual controller “boxes” and their electronics. This is edge computing, where IT cloud infrastructure is located close to the machines on the shop floor. This makes it possible to meet the stringent requirements of very low latency and short control cycles. In this paper we assume that the edge computing infrastructure is located on premise, within the factory or plant. This enables a paradigm shift, interconnecting a large number of vPLCs on an edge node to improve operational efficiency. Although vPLCs have been only adopted sporadically today, this architecture could gain importance with the advent of edge computing. In data centers, virtualization of databases and applications on powerful servers have long been state of the art.

Virtualization is the ability to separate logical functions (software) from the physical device and to run on commodity hardware. Virtualization has reduced costs, increased flexibility and scalability, and improved reliability and performance in IT. Over the past decade, we have seen many vendors supporting their supervisory control and data acquisition (SCADA) and distributed control system (DCS) platforms within a virtualized environment, leading to more virtualization in OT environments.

More recently, we have seen a number of DCS vendors deploy virtualized controller CPUs in situations when they need to increase performance beyond the capabilities of their current line of controllers and to reduce the cost of controllers for strategic accounts. These virtualized DCS controllers run on commodity IT hardware such as Windows or Linux servers.

## **1 THE INTEGRATED, EDGE-BASED ARCHITECTURE**

---

The following situation is conceivable: discrete PLCs are removed from the shop floor and their control functions are hosted in an *edge data center* in the form of vPLCs, with suitable computing capacities and network connection to the automation system. Servers can already process hundreds of vPLCs at the same time, but standard IT servers with virtualization and shared resources for network and storage cannot adequately address OT industry requirements for reliability and real-time behaviors. For this, industrial edge servers are required that are able to run real-time workloads. Only the I/O stays local to the machines, sensors, actuators and drives. Some examples of industrial workloads including real-time control are shown in Figure 1-1. *The manufacturing industry is adopting Edge Computing for IT/OT convergence*

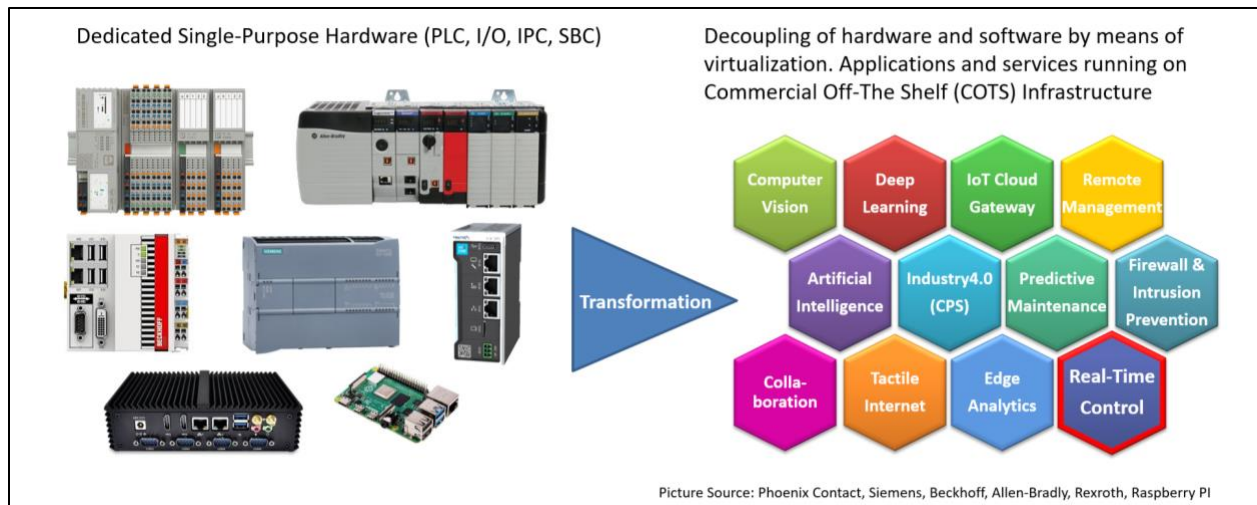


Figure 1-1. *The manufacturing industry is adopting Edge Computing for IT/OT convergence*

With no functional changes, an integrated architecture dramatically decreases capital and operational expenditures compared to a decentralized architecture based on individual PLCs, since a large number of virtualized PLCs can be hosted on a single server.

A clear difference between vPLCs and conventional PLCs can be seen in the flexibility and expandability. By virtualizing control functions and running them in the edge data center, interactions between virtual controllers become simpler. The communication among virtual controllers can be implemented by functional calls within a single server, which increases reliability and scalability compared with the traditional communication between physically separated PLCs. It dramatically facilitates update and re-design of the production line. With virtual control functions running in the data center, it behaves as a “digital twin” of the production line, which helps to simulate and predict the behavior of the physical counterpart.

The possibility of accessing the data from the field level in the edge data center means that controls and data analytics can be carried out in real-time, which is ideal for diagnostics, maintenance, optimization and intelligent reactions to changes in the automation system. The big data analysis doesn’t run “on” but “parallel to” the controls on the same edge servers. Therefore, modern AI and machine learning algorithms could be applied here without interfering with the existing control process. Since the control functions are to be hosted on the same edge infrastructure, feedback loops from data analytics to the control can also be implemented here, thus opening up new optimization options.

## 2 IMPLEMENTATION CONSIDERATIONS

Why is this architecture not (yet) prevalent? Surveys have identified four arguments:

- 
1. *The current architecture is tried and tested:* The cost advantages, flexibility and optimization are desirable, but the gains are not worth it. “Never touch a running system” except when there are significant gains in capital and operational expenditures.
  2. *Service-level agreement and liability:* The factory and plant operators bought a “closed” solution from the vendor or system integrator. This often has closed interfaces, so the machine cannot be adapted to a different control architecture. If the machine is not working as expected, the vendor or system integrator must fix it.
  3. *Technological risks:* The reliability and determinism of integrated server platforms are not trustworthy enough to outsource critical control functions to them. The response time of controllers in an edge data center may also be unreliable due to the network.
  4. *Organizational hurdles:* To implement an integrated platform of this kind, the control specialists need new skills. The distribution of competencies in the company is often incompatible with an integrated architecture.

The advantages of virtualization of control at the edge must be judged for each application. Edge computing infrastructure is particularly worthwhile for applications that place high demands on flexible production processes and reactive process changes as part of an Industry 4.0 strategy (batch size one).

The technological risks are quite challenging, but some recent developments show what is already possible here:

The advantages of virtualizing the control at the industrial edge must be judged for each application. In real-time data acquisition and edge analytics, advantages can be found in applications that rely more on flexible production processes and reactive process changes as part of an Industry 4.0 strategy (batch size one). Because of its flexibility, an edge computing infrastructure is particularly worthwhile for applications that place high demands on dynamic changes in automation. The technological risks are quite challenging, but some recent developments show what is already possible here:

- Real-time operation systems and hypervisor solutions already offer mechanisms for guaranteed robust partitioning of resources such as CPU cores and cache for current multicore CPUs, so that instead of virtualization and the associated runtime fluctuations, real-time performance like 'bare metal' can be achieved. The more cores, the more real-time applications can be run simultaneously and independently of one another.
- Hardware-supported network virtualization of the local Ethernet interface(s) enables several applications to use the network resources on the same server independently of one another and required bandwidth in the network is available in real time. Time Sensitive Networking (TSN) and Deterministic IP provide the mechanisms in the network to transport real-time data from different applications independently and without interference with guaranteed latency through an Ethernet network.

- 
- For various fieldbus protocols such as Ethercat and Profinet, specifications are already available to define 'tunneling' via TSN. This enables I/O at the field level to communicate with the vPLCs in the edge data center through one network, as if they were directly connected via the fieldbus.
  - There is a strong trend towards manufacturer-independent interfaces for application and management. The current specification work in the Motion Working Group of the OPC Foundation aims to ensure vendor-independent interoperability between controls and motion control devices (drives, I/O) based on OPC UA Publish/Subscribe. Similar standards for manufacturer-independent interoperability based on OPC UA and TSN have already been adopted by other industrial consortia and associations (e.g. Euromap and OPAF). Manufacturer-independent management interfaces may also play a role.

With these mechanisms, technologies and standards, the integrated and edge-based architecture can be implemented, as it already is in various test beds and demonstrators.

Nothing changes in the model of the control with regard to programming and runtime behavior: The IEC 61131 programming model can be used unchanged, even if the resulting control application is implemented as a virtualized PLC and only communicates with the field level via fieldbus interfaces. Meanwhile individual developments are already aimed at expanding the integrated architecture in accordance with the requirements of IEC/EN 61499 for distributed controls, e.g. <https://www.eclipse.org/4diac/> and <https://www.nxtcontrol.com/en/>.

Finally, the organizational challenges are non-negligible. The configuration, commissioning and maintenance of an edge platform for hosting of controls is a completely different task from programming and operation of a PLC. Therefore, the lifecycle of the control applications remains the domain of the experts for these applications as before, while the IT department takes care of the installation and maintenance of the edge servers and the hosting infrastructure running on them. The essential interfaces, which ensure robustness and determinism in accordance with the requirements of the automation industry, must be contained in a correspondingly certified edge computing platform product, because this competence is usually neither with the experts for controls nor in an IT department.

But this does not change anything in the basic model. Figure 2-1 shows the potential architecture of a hypothetical automation platform.

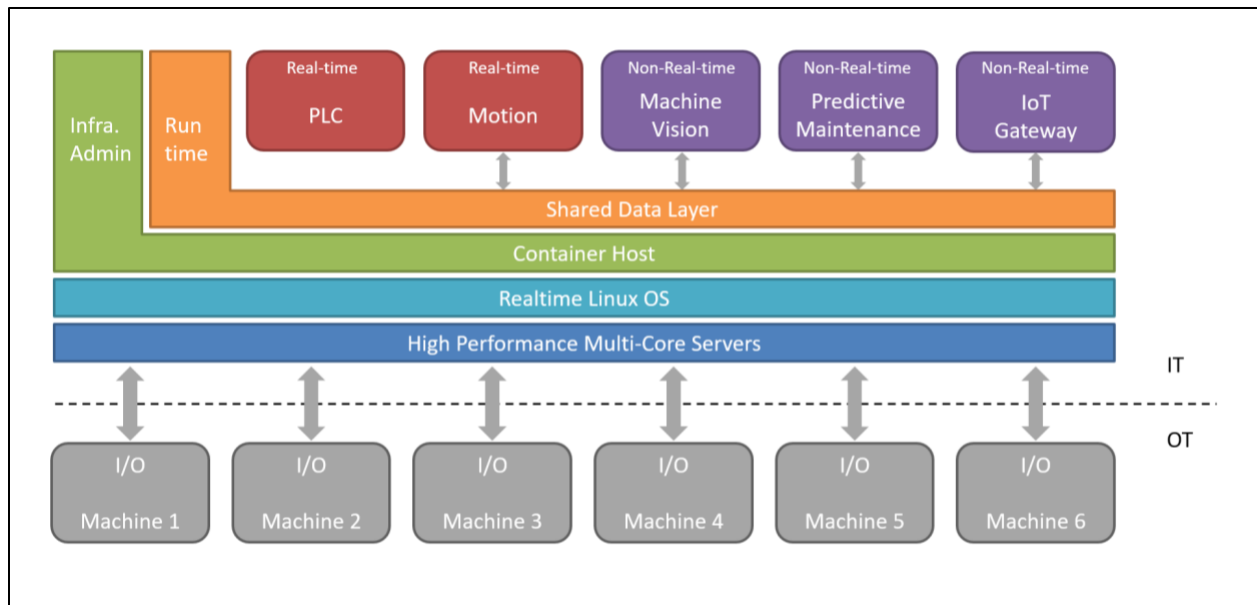


Figure 2-1. Automation platform architecture

### 3 A LOOK INTO THE FUTURE

A consistently integrated edge-based architecture for controls in automation is currently not state of the art. The current architecture, in which most controllers are implemented directly in its application as a hardware PLC, is well established. However, if there is a need for more flexibility, the methods and technologies for integrating the virtualized controls in an edge computing architecture with edge nodes and edge data centers can offer great benefits.

Certainly, there are challenges ahead for full PLC virtualization to become a reality. For example, there are fundamental differences between the deterministic nature of PLCs and the undeterministic, performance-focused nature of other traditional cloud services, e.g. office applications. Full PLC virtualization is unlikely to occur without one or more vendors getting involved in this technology shift. The vendors that create this reality would have greater market influence as the “VMWare of OT.”

### 4 BUSINESS CONSIDERATIONS

Although using virtual PLCs in industrial control opens up an opportunity to integrate all the advancements in ICT technologies developed recently, the adoption rate in the manufacturing industry can be improved. We recommend vPLCs be deployed for greenfield deployment. To get the maximum benefit, they should be integrated with actuators and sensors on the production line, and make the controller-with-vPLCs be part of the Smart Factory supporting subsystems.

---

Huawei is currently performing proof-of-concept validation in the process industry and in factory automation. Some promising results can be found in a paper we published together with Festo: <http://opendl.ifip-tc6.org/db/conf/cnsm/cnsm2019/1570581161.pdf>

## 5 RELATED PRESENTATIONS IN IIC

---

“Deterministic IP enabling virtualized PLC at the Edge”, IIC Innovation TG, July 28<sup>th</sup> 2020  
<https://engage.iiconsortium.org/wg/AllMembers/document/previewpdf/22201>

## 6 ABOUT THE AUTHORS

---

*The views expressed in this IIC Technical Brief are the contributing authors' views and do not necessarily represent the views of their respective employers nor those of the Industry IoT Consortium.*



**David Lou** graduated as Ph.D. in Electronic Engineering at Ghent University in 2005. In the same year he joined the Alcatel-Lucent Bell Labs as an Innovation Researcher. He had a leading involvement and management role in several European and national research projects (Giant, Smart Touch, Metaverse1, Mistra, Shift-TV, etc.), and standardization bodies (MPEG). In 2016 he joined Huawei Technologies as a Chief Researcher based in Munich, Germany. He is responsible for defining the research strategy, steering disruptive network innovation and coordinating collaboration with industrial and academic partners. He is also leading the standardization activities in various SDOs (e.g. ITU-T, IETF, ETSI, etc.) His interests mainly covers IoT/IIoT/I4.0, next generation industrial networking architecture, deterministic communication, network privacy, video streaming and transportation, and immersive communication. He is the co-chair of the IIC Networking Task Group and has been actively involved in relevant industrial development activities. He has (co-) authored more than 30 scientific publications and white papers. He has been granted with more than 20 patents.



**Ulrich Graf** holds a diploma degree in Electrical and Electronic Engineering (1994). Before joining Huawei in 2011 he worked almost 17 years in the telecommunication industry in various sales, product and R&D departments at Siemens and Nokia where he could extend his business & professional competence in the areas of Engineering, Product- and Program-Management for enterprise, fixed and mobile networks. Since joining Huawei Ulrich was established in the European Solution Management Team serving solutions like Policy/Traffic Management, Cell Congestion Control, Telco Transformation (SDN/NFV/MANO), IoT and



---

Industry4.0. In 2019 he joined the Applied Network Technology Laboratory in the Huawei German Research Center in Munich working on new communication and networking technologies targeting different industry verticals.



**Dr. Mitch Tseng** is a well-known veteran in the international standards community for communications. He had been leading the development of various telecommunication standards in TIA, 3GPP2, 3GPP, and helped established oneM2M prior to joining the Industry IoT Consortium (IIC). His latest passion is to promote technologies such as Edge Computing, Digital Twin, AR/XR, Li-Fi and AI for Industrial IoT (IIoT) Services. He is currently the Chair of the Innovation Task Group, the Edge Computing Task Group and the Testbed Council in IIC. In addition to the IIC work, Mitch is the Rapporteur of ISO TC204 WG16 (ITS Communication Protocols).

As a self-proclaimed IIoT Evangelist, he has been visiting end-users in different trades to listen and learn the Do's and Don'ts about the practices from them, so as to help refine the requirements for IIoT Services. Leveraging IIC as a vehicle, Dr. Tseng is also engaging in connecting members and outside partners to build a healthy IIoT micro-ecosystem. With all the meetings among interested parties in Asia, Europe, and Americas, Mitch's next milestone, literally, is to become a Four-Million-Miler of American Airline's Advantage Program.

Copyright © 2021, Industry IoT Consortium, a program of the Object Management Group ©. All copying, distribution and use are subject to the limited License, Permission, Disclaimer and other terms stated in the Industry IoT Consortium Use of Information – Terms, Conditions & Notices, as posted at [https://www.iiconsortium.org/legal/index.htm#use\\_info](https://www.iiconsortium.org/legal/index.htm#use_info). If you do not accept these Terms, you are not permitted to use the document.