Yokogawa Cloud Architecture for Smart Manufacturing

M K Naveen Kashyap [*1]

Digital transformation (DX) is the application of digital technologies and the power of information to an organization's people, processes, and systems with the aim of radically reorienting its business performance. Smart manufacturing is a form of DX applied to the manufacturing side of an enterprise. The goal of many smart manufacturing initiatives is the realization of autonomous operations. Key among factors that define the success of DX and smart manufacturing is digital enterprise architecture, which captures the tech stack and its interplay with existing systems and business processes.

DX architecture captures the entire value creation process, from data capture to value delivery, by way of various software applications and augmented offerings. This paper describes the essence of Yokogawa's DX architecture in terms of six layers, starting from the "plant floor," which combines existing production assets and associated IT assets. The topmost layer comprises augmented offerings that allow services such as consulting, engineering, system integration, and support to convey the created value to the customer. The four layers in between—edge, computing infrastructure, platform, and applications—represent the key technological elements that make DX a potential reality.

INTRODUCTION

Digital transformation (DX) is the application of digital technologies and the power of information to an organization's people, processes, and systems with the aim of radically reorienting its business performance. Fundamentally, this involves using data from various sources along with algorithms to make insightful and well-informed decisions. Whereas operational technology (OT) reflects all production assets, information technology (IT) reflects the various systems governing enterprise business processes. New insights are derived by combining data from both streams. For most existing enterprises, it is a great challenge to implement transformational projects without taking into account existing hardware, software, and business processes as well as the work force.

The Smart Manufacturing Leadership Coalition (SMLC) defines smart manufacturing as "the intensified application of advanced intelligence systems to enable rapid manufacturing of new products, dynamic response to product demand, and real-time optimization of manufacturing production and supply-chain networks."(1)

The ultimate goal of many smart manufacturing initiatives is the transformation from industrial automation to autonomous operation and industrial autonomy. Yokogawa has termed this transformation "Industrial Automation to Industrial Autonomy (IA2IA)."(2)

Key among the factors that define the success of DX and smart manufacturing initiatives is the digital enterprise architecture, which captures the tech stack and its interplay with existing systems and business processes. At the heart of this architecture is a platform stack that provides various reusable services and joins together the various components of the solution. The platform stack is designed to be infrastructure/cloud-agnostic, providing application enablement and facilitating quick and effective delivery of software as a service to our customers.

DX architecture thus captures the entire value-creation process, from data capture to value delivery, by way of various software applications and their augmented offerings. Figure 1 shows an overview of a DX architecture in terms of the six layers, starting from the "plant floor," which combines existing production assets and associated IT assets. The topmost layer consists of augmented offerings that encompass services like consulting, engineering, system integration, and support in order to convey created value to the customer. The four layers in-between—edge, infrastructure (IaaS), platform (PaaS), and applications (Apps)—represent the key technological elements that make DX a potential reality. Note that this layering differs from the ISA95 pyramid or Purdue pyramid.

*1 Digital Innovation & Services, Digital Platform Center, Yokogawa Electric International Pte. Ltd.

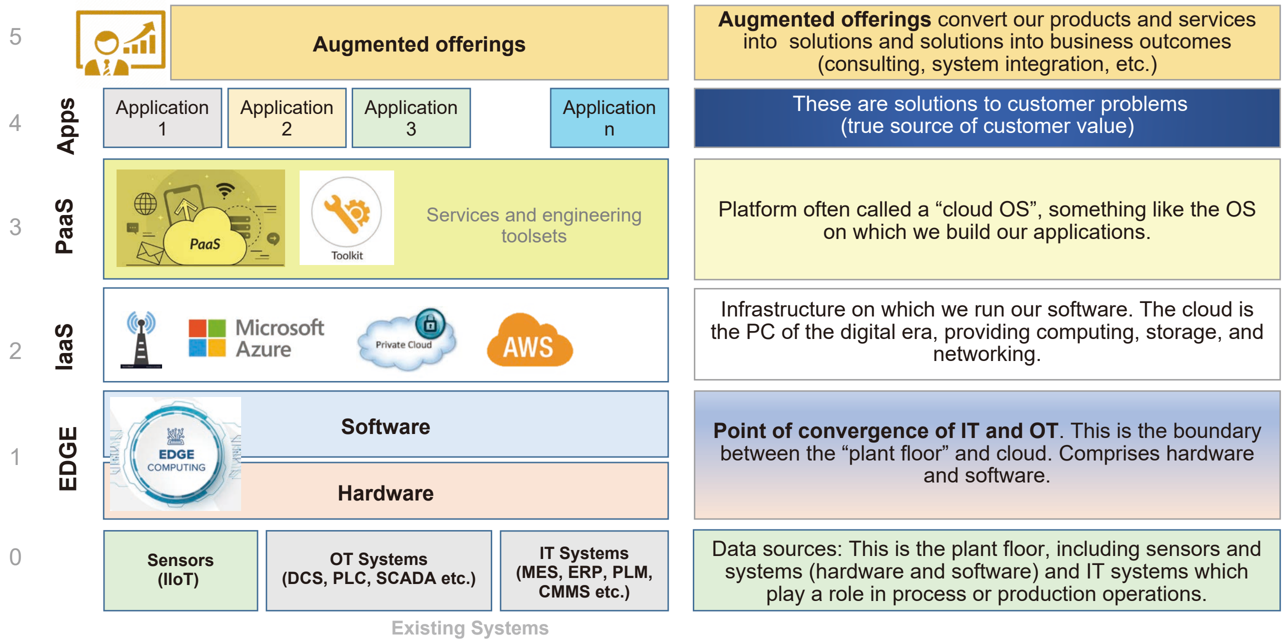

Note: The terms "digital architecture," "digital enterprise architecture," and "DX architecture" and Yokogawa's DX architecture for smart manufacturing, Yokogawa Cloud Architecture, are used interchangeably in this document.

Figure 1 The six layers of a digital transformation (DX) architecture

YOKOGAWA'S DX ARCHITECTURE

Yokogawa's DX architecture follows a layered approach to closely match the six levels of a typical DX strategy, as shown in Figure 1. The structure of Yokogawa's cloud is shown in Figure 2. The topmost layer, augmented offerings, is not shown because it is more concerned with business processes related to DX offerings that define the user experience.

Each of the six layers in described in the following sections.

Yokogawa positions itself as a software as a service (SaaS) provider in the digital economy, in that it licenses apps and services to the end user, rather than the platform itself.

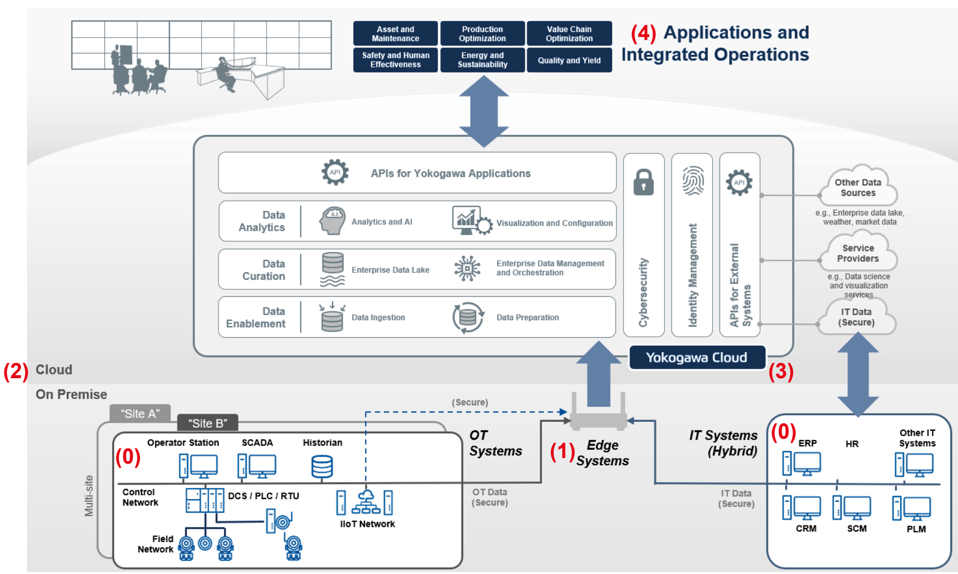

Figure 2 Yokogawa Cloud – an architecture for smart manufacturing

LEVEL 0: OT AND IT

The OT part consists of systems primarily on three types of networks:

(1) Various systems on the control network, such as distributed control systems, programmable logic controllers, remote terminal units, and data acquisition systems, along with a wide range of field equipment for communicating with these protocols, including 4–20 mA, HART, FF-H1, ISA100, and MODBUS.

(2) SCADA, plant historians and other plant information management systems on the plant information network.

(3) Wireless sensors that measure and transmit information either to edge equipment (preferred) or directly to a sensor cloud, using protocols such as ISA100, LoRaWAN[3], Sigfox, and other wireless protocols.

IT systems typically include MES, ERP, CRM, PLM, HRMS, and other systems. Most provide application programming interfaces (APIs) based on standard industry protocols that facilitate data exchange between systems. The configuration is simplified in Figure 2.

LEVEL 1: EDGE

"Edge," indicating local data processing and computing, is a confluence of enterprise IT and OT systems. Edge computing is computing performed at or near the data source, instead of performing all the work on the cloud or in a data center. The edge is often considered the bridge between the old (OT) and new (IT) paradigms and between the "south side" (the plant floor and OT) and the "north side" (the cloud); an edge controller is often the secure endpoint for all plant floor systems. Figure 3 illustrates the idea of a unified edge.

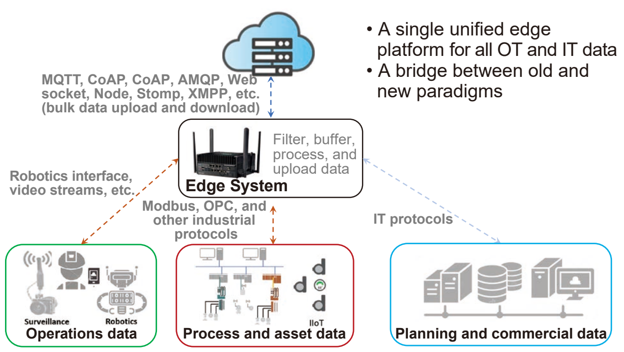

Figure 3 Yokogawa Edge – a unified bridge between multiple paradigms: OT and IT, old and new

Ideally, the edge should provide an endpoint not only for process data, but also for other data relevant to manufacturing, such as operations data (e.g., shift information and robotics data) and IT data (e.g., planning and scheduling).

One way in which an edge system differs from present-day distributed systems is the myriad computing possibilities that it brings to the use of unstructured data and artificial intelligence (AI).

Key considerations for edge design include a clean data ingestion strategy from existing or legacy plant systems to leverage the power of data analytics and cloud computing, without interfering in their routine functions. The edge consists of a combination of hardware and software. The sizing and deployment of the edge also depend on factors such as latency, bandwidth, privacy and security concerns, and most importantly, the nature of the application. An edge gateway or edge server has two roles. The first is to regulate the flow of information from the OT and IT floor to the data center or cloud (and vice versa), providing storage or buffering of data wherever applicable. The second role is to provide some level of intelligence to perform data filtering and execution of business logic closer to plants or to machine-learning or transfer-learning algorithms, as required. The edge requirements often depend on the application. The choice of hardware ranges from commercial off-the-shelf (COTS) systems to proprietary Yokogawa hardware.

The edge acquires OT data from process systems or process historians using interfaces conforming to open industry standards like classic OPC and OPC-UA; native interfaces that support a host of other industry standard protocols like MODBUS, Foundation Fieldbus (FF-H1), and Profibus, as applicable; or other web services where available, regardless of the physical network layer. In the case of industrial internet of things (IIoT) applications, edge nodes also serve as message queuing telemetry transport (MQTT) brokers (or constrained application protocol (CoAP), extensible message and presence protocol (XMPP), restful HTTP, etc.) to gather data from edge devices that act as clients. In addition, these systems are capable of bulk upload of data gathered from various systems. There are also provisions for using bridge devices for protocol conversion (both at the physical and informational level) to link OT systems with the edge. The standard security practice is to deploy the edge in a DMZ configuration, although other options for deploying edge servers with information diodes on the network are available and can be defined on a case-by-case basis. Edge gateways can also be configured and managed from an external network like the cloud, based on organizational security policies. In some configurations, edge functions can be partly located at customer data centers as required and have the flexibility of pushing some functionality from the cloud to the edge as necessary ("scalable edge").

An enterprise may consist of multiple sites, and thus have multiple OT systems and networks with the configuration described above. The edge as part of the hyperlocal/hybrid cloud strategy focuses on local data processing, data management to and from the cloud, control, autonomous operations, and resilience (in case of disconnection from the cloud). The central cloud connects to multiple site-level edges and focuses more on data aggregation (OT and IT) from multiple sites (as illustrated in Figure 4), AI and machine learning (ML) applications (learn from one site and deploy everywhere), and integration of the enterprise with third-party cloud systems (e.g., supply chain–partner integration).

Figure 4 illustrates the role of the edge from the perspective of site and enterprise operations.

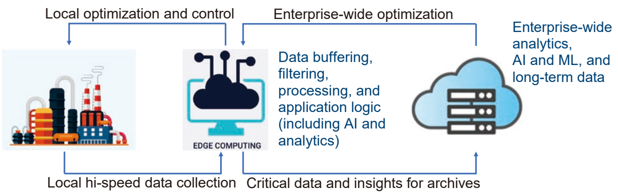

Figure 4 Edge and Cloud – connecting local and enterprise operations

Edge services can also be extended to IT systems when bandwidth, real-time processing, latency, or synchronization of IT and OT data near the source are key application requirements. The edge system then becomes a unified endpoint for both OT and IT data, thereby eliminating the need for IT systems like ERP or PLM to establish their own data paths to the cloud.

The edge strategy would then vary between customers and their plants, as well as the application. Yokogawa adopts a three-tier edge strategy depending on the application.

- The first tier is the "Edge Gateway," which is focused on providing secure transport of data from the plant floor to the Yokogawa Cloud, but with restrictive OT protocol support. This software could run on industrial computers and Yokogawa's e-RT3 hardware.
- The second tier is the "Light Edge," which in addition to providing all functionalities of the basic tier, also provides connectivity across a broader range of OT protocols along with buffering, filtering, payload transformations, and additionally some IT data enablement.
- The third tier is the "Comprehensive Edge," which in addition to all first- and second-tier services, also provides application enablement, that is, the running of applications on top of the edge (including AI apps), application orchestration, device management, support for robotics, and remote engineering from the cloud. Yokogawa is currently standardizing its edge software strategy, called "CI Edge," combining components of the Collaborative Information server with those components for video and image analytics provided by amnimo and Yokogawa's cloud platform explained in the section "Level 3: The Digital Platform."
- The "end state vision" of the edge supports ingestion of both structured data (from sensors, process data, etc.) with unstructured data (video, images, files, etc.). This would ensure that our edge is able to process and bridge all kinds of operations data (OT and IT) as described in Figure 3, between the plant floor and the cloud. The ability to execute AI applications that combine one or more of these data will unlock new value for customers through applications like virtual/smart workers, integrated remote operations, robotic applications, etc., thereby taking our customers a step closer to realizing industrial autonomy. By applying distributed machine learning (DML), it becomes possible to distribute the AI workload seamlessly between the cloud and edge, enabling a new paradigm of smart manufacturing.

LEVEL 2: CLOUD INFRASTRUCTURE (IAAS)

An edge server pushes data to the "cloud," meaning one or more private data centers (on- or off-premises) or to a public cloud infrastructure like Microsoft Azure, Amazon Web Services (AWS), or Google Cloud Provider (GCP). Yokogawa adopts a cloud-agnostic strategy (i.e., support for multiple cloud provider infrastructures).

Y-AWS and Y-Azure are cloud/infrastructure services on the Azure or AWS public clouds that are secured, managed and operated by Yokogawa, as part of our fully managed cloud software services offerings to end users. These are often referred to as private clouds on the public cloud infrastructure. The choice of cloud provider depends on the footprint of the respective operator in each region, along with local regulations regarding data sovereignty, localization, and data residency. Yokogawa offers its digital services on such instances or instances of the customer cloud on Azure or AWS or on-premises infrastructure. Hybrid cloud deployments, where data remain on the customer cloud, but applications reside on one or more public clouds, are also supported. Hybrid and on-premises configurations vary from case to case and thus might or might not support all applications or features. They are thus supported on a case-by-case basis.

LEVEL 3: THE DIGITAL PLATFORM (PAAS)

The digital platform is regarded as the operating system for the cloud or as middleware for digital apps. The platform provides many reusable services to accelerate application development and delivery of digital services and incentivizes collaboration and sharing of resources between them (applications). The objective of the platform is to provide a core of reusable services around which the regions and businesses can build solutions in a quick, simple, and efficient way and deliver them to customers using the cloud. The platform also promotes the sharing of data and other functions between applications by way of an APIs. The vision of the platform is to eventually cut the time and cost required for new applications and solution development by half.

Four key design principles of the platform—connectivity, flow, convergence and gravity—enable the above sequence (Figure 5). Connectivity is a measure of how easily various applications and systems securely connect to the platform, whereas flow is how effectively connected applications seamlessly exchange information through standard interfaces. Convergence is the ability of the system to integrate, combine, and sanitize different data sources into a useful form ready to be consumed by other applications. Gravity is the ability of the connected system to add value by attracting more applications and services into the ecosystem.

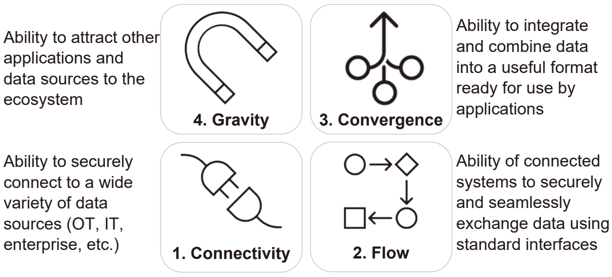

Figure 5 Core principles of platform design

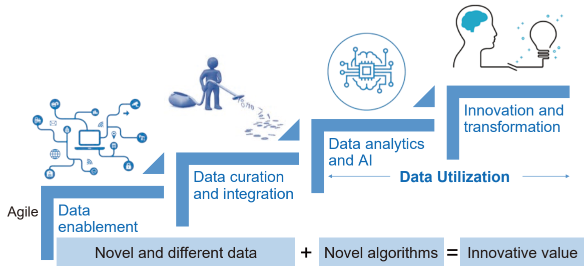

Figure 6 Data-centric principles for platform design

Yokogawa's digital innovation platform goes beyond just an IIoT platform that supports IIoT applications, covering the entire spectrum of digital applications and is thus positioned as a true DX platform. Given our positioning as a SaaS provider, the platform tooling is available only to Yokogawa applications built upon the platform, and platform engineering functions are available only to Yokogawa engineers. Our end users directly experience the applications hosted on the platform and not the platform itself, which is a hidden enabler. The design of the platform follows a data-centric architecture that manages data along its entire lifecycle from ingress (entry) to egress (exit), and is architected around the established sequence of activities and processes around data, namely, *data enablement* (getting data into the system), *data curation* (transforming, storing, and organizing data so acquired), and *data utilization* (consuming the raw, acquired or process data, in order to perform analytics and ML to derive insights), culminating in innovation (Figure 6).

The platform is organized into five layers (3-1 through 3-5 as illustrated in Figure 7), each representing a set of tools and services available to applications built on the platform. Common services (3-1 in Figure 7) include cybersecurity and identity management services that provide secure access to data. The layered architecture of the platform maps to the key phases of data management: data enablement (3-2 in Figure 7), data curation and processing (3-3 in Figure 7), and data utilization illustrated as analytics (3-4 in Figure 7) by virtue of being the most compelling case of data utilization. It also includes an API (3-5 in Figure 7) through which the applications consume the services provided by the platform. This principle is illustrated in Figure 7.

Each of the five layers of services mentioned above and illustrated in Figure 7 is explained in detail in the following sections.

3-1 Common services

Identity and authorization management, and cybersecurity related services are grouped together and often referred to as common services, as explained in detail in the following sections:

(a) **Identity and authorization management** for user management, authorization, single sign-on, and other user or role privileges that enable secure access to platform

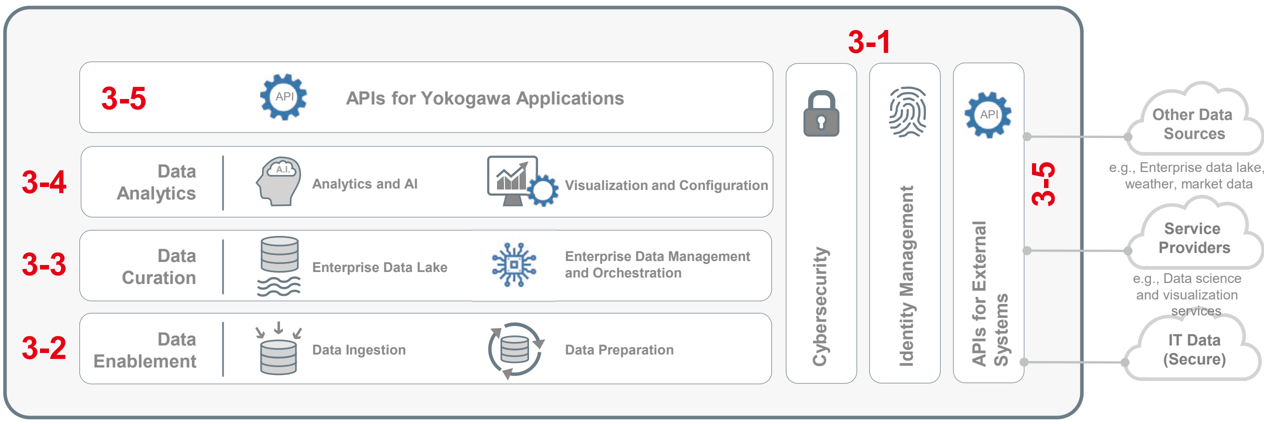

Figure 7 Yokogawa cloud platform

resources. The platform is also capable of linking to external identity management systems used by the customer.

(b) **Information security management** encompasses data security, encryption, and segmentation throughout the lifecycle, including data in transit and at rest as well as for APIs that enable information exchange.

3-2 Data enablement services

These services focus on bringing data from multiple sources into the platform and performing basic janitorial services (primary data clean-up). Among others, it includes the following services:

(a) **Data ingestion:** These services include the ability of the system to move data from multiple sources (both OT and IT) into the platform. Data enter the system in various formats, either in real time (synchronously) or in batches (asynchronously). Data types include process data, alarms, alerts, files, and objects. They also include functions for device management and on-boarding from an IIoT perspective. As a DX platform (compared with most IIoT platforms), the data goes beyond sensor data to include other process information.

(b) **Data preparation:** These functions permit the user to cleanse the data, filter and normalize them as required by way of tools so that the ingested data can be converted to a form that can be used by the next layers of the platform and applications.

3-3 Data curation services

These tools help organize, describe, clean, enhance, and preserve data in a form that can be seamlessly reused by multiple applications. The scope of curation includes both data and metadata that define the data context. The processed data can be stored in a database or in a data lake. In addition to providing storage, data curation integrates data into structures, models, and repositories that are more valuable than raw data. These tools embody Yokogawa's strong understanding of the domain and operation contexts across the levels of the ISA95 pyramid, including our proven data management processes and practices from our control systems, process historians, and our security management services. The data curation tools enable app developers to quickly build applications around ISA95 and other standards-based information models.

The platform design ensures that customer data is securely managed and governed, providing segmentation and isolation from other data. The key services at the data curation layer include the Enterprise Data Lake services and the data management/ orchestration services, which are described below.

(a) **Enterprise data lake services** provide centralized storage and management services for structured and unstructured data at an elastic scale. When storage is in a data lake, the philosophy is to store all raw data along with processed data. Data storage can also be in a database instead of the data lake in itself. It is important to note that the platform provides the services needed to manage data flows in and out of data lakes, rather than providing a full-scale enterprise data lake itself. This is because the enterprise data lake implementations are specific to each organization and can seldom be standardized. Organizations in most cases have enterprise data lake implementations that cover a larger scope of the enterprise operations and expect to be able to access the same data lake functions, rather than creating another data lake, given that data lakes are expensive to create, operate and maintain. When an external data lake exists, the platform can utilize data lake services instead of duplicating them.

(b) **Enterprise data management and orchestration services** span the entire set of tools and frameworks that enable cataloging, organizing, managing, and processing data for delivery to specific analytic applications. These services help analytics applications automate the different steps of data pipeline development from source to consumption. These include extract-transform-load operations, data transformations, and data modeling to provide a fully abstracted but unified environment for delivering large volumes of data for individual analytics applications in the form they expect.

3-4 Data Analytics

Data analytics is one of the most important services provided as part of the "Data Utilization Phase" of the data life cycle explained in Figure 6.

Processed data is made available for various applications, including data science applications. This layer includes low-level code tools, graphical logic builders, exploratory analytics, a sand-boxed environment for experimental analysis and machine learning, dashboarding tools, and engineering tools comprising the business logic and user interface tools for powerful visualization. The key components are explained in the following sections.

(a) **Analytics services:** These services include basic tooling and support for off-the-shelf analytics and algorithms for time-series and batch data. The platform also supports containerization that allows packaging of software into standard units (with resource isolation) for development and deployment. Support for AI and ML algorithms is also grouped under analytics, with the option to integrate third-party AI frameworks or offerings to enhance the utility of the application. This unique mechanism of welcoming third-party AI code based on standard frameworks like TensorFlow and Pytorch is called "Bring Your Own Machine Learning" (BYOML). In addition, the services include a strong Extract-Transform-Load (ETL) pipeline service which allows the user to utilize the data insights from the platform from third-party business intelligence (BI) tools like Microsoft Power BI and Tableau. This function is called "Bring Your Own BI" (BYOBI).

(b) **Visualization services:** The platform also provides a visualization framework, built using a standard design language that aids creation of user interface elements like dashboards and displays. This layer comprises design-time and run-time frameworks that respectively support application creation or engineering and application usage. The platform also provides hooks to popular third-party visualization tools so that the application's user experience is consistent with any analytical packages already in use. Visualization is also available on mobile devices by way of an application. The vision is to meet different kinds of visualization needs of different kinds of apps like BI dashboards, monitoring UIs, process displays, etc. starting from BI dashboards.

3-5 Application programming interface (API)

APIs are interfaces that allow various application components to communicate and exchange data as appropriate with platform components or with other applications. Of the myriad existing methods and architectures for designing such APIs, REST is popular and uses HTTP as the protocol for facilitating communication between software or its components, like web browsers and web servers. Beneath the APIs are (typically) microservices, a collection of loosely coupled, modular web services, each performing specific functions.

Every significant platform function is exposed as a REST API through an API gateway for exclusive use by Yokogawa apps. The APIs are secured by secret keys to prevent abuse and misuse. A second set of APIs allows the platform to exchange data with external systems or platforms. These include external data sources on other clouds, including enterprise data lakes, partner systems, or analytics service providers.

LEVEL 4: APPLICATIONS

Applications or solutions are software programs built to solve specific business problems, such as asset management, production optimization, or health and safety. Applications realize the vision of smart manufacturing and DX.

Applications are built using services such as data, enterprise data management and orchestration, and the logic builder and visualization available from the platform. Platform plumbing encourages applications to connect with each other and securely exchange or reuse data as needed. Once available on the platform, data can be reused by other applications, thereby eliminating data duplication throughout the lifecycle. Standard applications are configured for deployment using logic and graphic builder functions. The application portfolio consists of standard and custom Yokogawa applications built for specific regions.

At the time of this publication, the platform services are exclusive to Yokogawa applications and more than fifty applications are on the roadmap. The platform has the ability to access (provide) data from (to) third-party platforms or applications via RESTful APIs.

LEVEL 5: AUGMENTED OFFERINGS

Augmented offerings consist of all activities forming the digital value vehicle for our customers. They are a collection of processes in which regions and frontlines play a dominant role in defining and sustaining the customer experience. Activities include DX consulting, application engineering, system integration, deployment, training, and post-sales support. Such processes are enabled by robust back-end systems managed by Yokogawa's Digital Strategy Headquarters.

KEY DIFFERENTIATORS OF OUR PLATFORM

Yokogawa's expertise in measurement, information and control, and its accumulated expertise in handling field, process and operations data (between Yokogawa and KBC) help us build platform tools more suited to process automation than many of the general-purpose IIoT platforms available in the market. The platform is tailored to *OT by design* as exemplified by our vision for:

– strong support for plant-side connectivity (more than 40 industrial protocols through the CI Edge, connectors to historians (OSI-PI, Exaquantum, Fast/Tools, etc.),
– progressive support for ISA95 related information models that support integration with MES systems,
– a process industry-focused application portfolio that goes beyond simple asset visualization, and
– dashboarding solutions provided by most IIoT offerings in the market.

The other features of the platform are also designed with smart manufacturing use cases in mind.

While the Yokogawa Cloud platform and solutions use infrastructure and services provided by hyper-scalers like AWS, Azure, etc., the underlying tech stack is not locked to any one of them. The platform stack is an integration of best-of-breed technologies and offerings, carefully chosen from a wide variety of sources. Yokogawa Cloud builds an additional layer on top of the basic services provided by hyper-scalers with the aim of progressively reducing the time and effort for application development.

Our cloud agnosticism ensures that the platform abstracts the applications and services from the underlying infrastructure so that apps built on our platform run seamlessly across different cloud infrastructure. Our portfolio of OT connectors, data- and function-sharing between applications, frictionless user experience across different applications of the platform, innovative approach to creating solutions by way of low code logic, and resource templates offers a high degree of control and support to quickly build and operate cloud-native applications.

Some of the differentiators include:

(1) The platform is cloud-agnostic (Azure, AWS, and on-premises platforms on a case-by-case basis).
This provides flexibility in combining infrastructures, platforms, and applications as a totally managed service bundle, or the choice of only applications (on an infrastructure of the customer's choice). The platform is a hidden enabler.
(2) Its vision supports the overall scope of DX, not just IIoT, meaning that we cater to a wide range of operational data other than sensor data, including relational data. Future

releases of the platform will focus on expanding its ability to ingest diverse sources of data including unstructured data like videos and images.

(3) It allows the integration of out-of-the-box time series analytics functions of the platform with any compatible third-party machine learning framework, a feature that we call BYOML. This function provides a run-time environment for models generated or tuned in an external environment (third-party system, as long as the model follows a set of standard frameworks).

(4) Strong edge capability (when available in future) characterized by a set of forty or more OT protocols, combined with an IT connectivity stack hosted on reliable Yokogawa hardware like e-RT3, and supporting execution of AI algorithms for image and video ingestion, based on amnimo's core technology makes it suitable for innovative IA applications (amnimo is Yokogawa's subsidiary for IIoT services).

(5) Seamless deployment of an application at the edge or on the Yokogawa Cloud provides deployment flexibility and a phased migration to the cloud.

CONCLUSION

DX reorients the business performance of our customers through the judicious application of digital technologies to all areas of their operations, including value-chain optimization, enterprise planning, asset management, and health, safety, and environment management. Our DX solutions include a variety of applications embedded with business wisdom built upon foundational technology blocks called the DX platform, typically (but not necessarily) on a cloud infrastructure. The DX platform provides a collection of common reusable services for data management, algorithmic execution, and visualization. By using a common platform, Yokogawa can create a variety of applications and integrated solutions with a consistent user experience, thereby accelerating and simplifying the value creation process. Business values unlocked by DX include improved effectiveness, efficiency, optimization, organization-wide collaboration, and progress towards industrial autonomy.

REFERENCES

(1) Stephanie S. Shipp, Nayanee Guputa, et al., "Emerging Global Trends in Advanced Manufacturing," IDA Paper P-4603, 2017, https://www.nist.gov/system/files/documents/2017/05/09/IDA-STPI-report-on-Global-Emerging-Trends-in-Adv-Mfr-P-4603_Final2-1.pdf (accessed April 20, 2021)

(2) Yokogawa Electric Corporation, IA2IA, https://www.yokogawa.com/solutions/solutions/ia2ia/ (accessed April 20, 2021)

(3) LoRaWAN® is a set of standards and specifications maintained by LoRa Alliance, https://lora-alliance.org/wp-content/uploads/2021/01/revised_lora_alliance_marks_and_logo_usage_policy_and_guidelines_06191.pdf (accessed April 20, 2021)