OPEN MANUFACTURING
PLATFORM

Insights Into Connecting Industrial IoT Assets

A publication by the Open Manufacturing Platform

Published	Version
December 07, 2020	1.0

Legal Disclaimers

“THESE MATERIALS ARE PROVIDED “AS IS.” The parties expressly disclaim any warranties (express, implied, or otherwise), including implied warranties of merchantability, non-infringement, fitness for a particular purpose, or title, related to the materials. The entire risk as to implementing or otherwise using the materials is assumed by the implementer and user. IN NO EVENT WILL THE PARTIES BE LIABLE TO ANY OTHER PARTY FOR LOST PROFITS OR ANY FORM OF INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS DELIVERABLE OR ITS GOVERNING AGREEMENT, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, AND WHETHER OR NOT THE OTHER MEMBER HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.”

Working Group Chairs, Authors, and Contributors

ACKNOWLEDGMENTS
This document is a work product of the Open Manufacturing Platform – IoT Connectivity Working Group, chaired by Sebastian Buckel (BMW Group) and co-chaired by Dr. Veit Hammerstingl (BMW Group).

AUTHORS
Open Manufacturing Platform / IoT Connectivity Working Group

CONTRIBUTORS

Attrey, Kapil	(Cognizant)
Birn, Lukas	(Capgemini)
Buckel, Sebastian	(BMW Group)
Hammerstingl, Veit, Dr.	(BMW Group)
Köpke, Alexander	(Microsoft)
Kramer, Michael	(ZF)
Krebs, Jeremy	(Microsoft)
McGrath, Daniel	(Cognizant)
Muth, Manfred	(Red Hat)
Reimann, Jens	(Red Hat)

Table of Contents

Table of Figures

1 Manufacturing at an Inflection Point

As the 4th industrial revolution unfolds, opportunities and threats present themselves to the entire manufacturing sector. Companies that embrace the change through experimentation and evolution will thrive while stagnant companies will struggle. Digitization is at the center of the revolution, and connectivity is paramount to digitizing all aspects of the manufacturing value chain.

In manufacturing, multiple challenges complicate the connection of sensors, actuators, and machines to a central data center. Lack of common standards and proprietary interfaces leads each engineer to solve similar problems, introducing inefficiencies and forcing the ascension of the same learning curve over and over. The long renewal cycles of shop floor equipment, software, and processes present gaps in modern technologies and a general avoidance of making significant institutional changes.

Each connectivity challenge will have a range of diverse constituents. Operational technology (OT) professionals are responsible for the commissioning, operation, and maintenance of shop floor equipment while information technology (IT) personnel look after overall data processing, the hardware and software infrastructure, and enterprise-wide IT strategy. General managers and logistics teams are typically aligned at a corporate level, coordinating strategy across a network of plants. Each of these functions will have roles spanning from operational hands-on to strategic and managerial. The unique demands of each part will require connectivity solutions to be forward-thinking and value-accretive while offering practical solutions implemented with minimal incremental investment. The overall goal is to transform data into information, information that can be leveraged by the stakeholders to create business value. Examples include: Optimizing existing business processes (e.g., Improving OEE (Overall Equipment Effectiveness)), building new business models, business-to-business data exchanges, and strategic partnering opportunities.

Raw data needs to be collected, aggregated across multiple data sources, consolidated, and analyzed to create information. Typically, data sources include shop floor machines, Supervisory Control and Data Acquisition-Systems (SCADA), Manufacturing Execution System (MES), Enterprise Resource Planning (ERP) systems, Distributed Control System (DCS) systems, as well as external supplier data. Historically these systems have been proprietary; containing locked down data formats and limited interoperability. Going forward, Industry 4.0 will necessitate the breaking down of these silos, and openness and interoperability will be essential to intelligent systems that persist into the future.

The Open Manufacturing Platform (OMP) was founded to harness the power of collaboration and openness, combining industry-leading companies and knowledge to solving challenges shared across the industry. Collectively innovating and working together creates practical solutions to real-world problems. Connectivity is the first working group because it is a cornerstone of unlocking additional scenarios, and this paper is the initial publication laying out an approach to solving connectivity challenges while providing a roadmap for future work.

2 Breadth and Depth

2.1 Industrial IoT Challenges

The Industrial Internet of Things (IIoT), as applied to manufacturing, has a different scale and time criticality than some other IoT platforms used in residential, fleets, or other non-industrial scenarios. Lower intensity use cases typically handle millions of devices that are all connected and sending relatively infrequent data. In general, relatively modest amounts of data accumulate in these systems, and messages (often less than 1KB) are sent at relatively infrequent intervals. For example, a smart temperature meter doesn’t need to publish the current temperature every second because it remains constant over such a short timeframe. Similarly, in car fleet management, it is not required to send data such as position or average speed every minute because the use case does not need it. This would be a waste of resources such as power, computing capacity, or data storage. Manufacturing, however, can be much larger in scale and scope.

Industrial IoT devices have critical real-time needs for repeatability and high availability. An example is an AI model that optimizes the parameters of a bending machine based on the current air temperature and humidity. Possible connection failures or high latencies can lead to stopped or interrupted processes or products with insufficient quality.

Manufacturing has a broad range of throughput requirements from low bandwidth for simple sensors using small packets to much higher bandwidth required for streaming data for video analytics, vibration sensors, or AR/VR visualization. A holistic connectivity solution needs to be considered up-front to address this complexity successfully, spanning from the individual devices on the shop floor up through edge gateways and servers to the central data center or cloud resources such as compute and storage.

2.2 Network Levels

Networks are mostly customized to their precise environment and the desired function, and therefore can be very complex. For the purposes of this discussion, we will consider a system with three logical levels:

- The cloud level describes globally available and scalable compute, storage, and other services running in a public cloud. It is remote from the production site.
- The edge level extends the cloud capabilities geographically closer to OT devices. It’s an on-premise network that possesses a path to the cloud as well as to the Sensor/Actuator level. Edge nodes that can run universal workloads are typically located at that level. These are often hosted on industrial PCs (IPC) and compute devices, virtual machines in a data center, or an on-premise Kubernetes cluster.

The production asset level is the network level where the OT assets are located, which capture the real-world processes (e.g., robots, cameras, machines) and react to them.

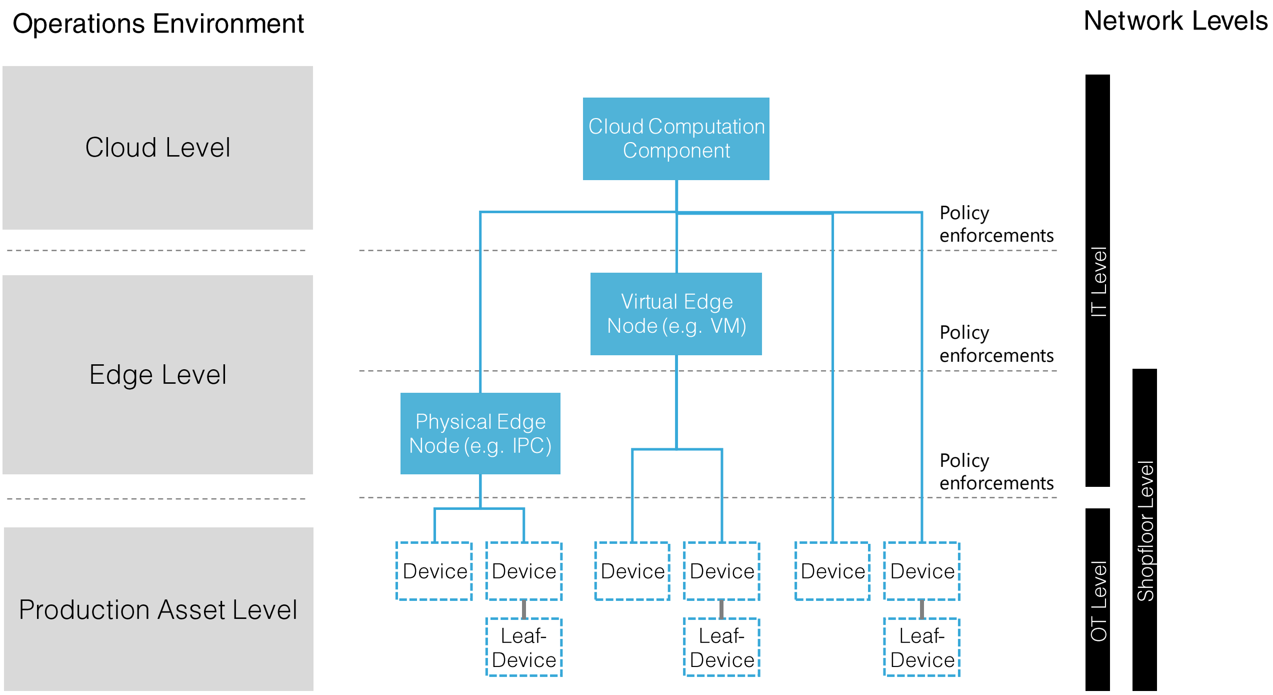

Device: e.g. PLC, Machinery, intelligent Sensor
Leaf Device: devices with limited communication capabilities e.g. basic sensor

Figure 1: Industrial IoT architectures typically comprise out of a production asset level, edge level and cloud level.

Equipment maintenance staffs operate on the shop floor level, which is comprised of the *production asset level* (OT) and physical edge nodes on the edge level, e.g., IPCs.

By deploying workloads on the edge level close to the data sources, the following benefits to IIoT arise:

- Security and data sovereignty - By processing data locally, data doesn’t need to be sent over the public Internet, potentially increasing security and enabling easier compliance with data sovereignty laws.
- Reduced latency - Local data processing increases response time and helps with reducing the latency for control loop tasks.
- Minimized bandwidth consumption - For high throughput use cases, like processing video feeds, it is often not feasible to send all data to the cloud for processing. Preprocessing, like filtering, can be done in this layer.
- Buffering - Data can be locally buffered if outages in later network layers occur.
- Equipment connection - Additional hardware components (e.g., cameras, sensors), even if not Internet ready, can be connected to the cloud.
- Process decoupling – Closed-loop operations are fully executed on the local device without dependencies on the broader network infrastructure

Depending on the given IT infrastructure, governance, or the kind of manufacturing, the production asset level could also be an aggregation of multiple network layers. For example, the Purdue network model was adopted to be a part of the ISA-95 (International Society of Automation) security standard, where an automation/control layer and additional layers are defined.

Securing the different layers of a given network is a complex business. Here are some example rules to potentially implement:

- Prevent a higher network level from accessing a lower level
- Explicitly grant access to the immediate one level lower network layer
- Prevent access to network layers that are not adjacent
- Ensure access only for well-known resources
- Setup an in-depth defense to prevent a breached layer from spreading to the others.

Scenarios where OT assets directly communicate with the cloud (e.g., via 4G/5G over the air networks) are often restricted due to security governance rules and require additional corporate government efforts. This scenario is also a viable architectural pattern for remote production sites, especially if an industrial PC is utilized.

3 Principles for a Successful Connectivity Solution

3.1 Leveraging the Cloud

Compute, storage, analytics, and other cloud services are secure, feature-rich, and highly scalable from essentially anywhere in the world. In today’s globally integrated networks of factories and supply chains, these features enable OT and IT managers to build, deploy, and grow solutions quickly and at low costs. These efficiencies are coupled with the tradeoff of control or freedom. Careful architectural planning (e.g., defining a business domain model) and the application of cloud principles, such as containerization, enable a technology manager to realize the benefits while minimizing the tradeoffs.

3.2 Building with Open Standards

Open standards have been used in the manufacturing sector for decades, providing functionality such as field and controller level communications and controller level operations. This approach has helped avoid some vendor lock-in, but there is room for increased proliferation and further leveraging of open standards. This is precisely the mission of the Open Manufacturing Platform.

3.3 Implementing a Platform Approach

As factories are both idiosyncratic and standardized along many different dimensions, any technology solution must account for the differences while leveraging the similarities. Devising a technology platform approach up-front enables rapid, consistent solution building and deployment while also allowing custom solutions that leverage a common core. Platforms are open for extension but closed for modification, thus providing scalability while minimizing breakage. They favor standard features and functionality over specifics tuned to one line or site. With stable, open, and well-documented APIs (Application Program Interface), provide a framework in which clients can implement their extensions with confidence and ensure interoperability. Additionally, security reviews of the core services can be centrally managed, and policies can be applied one-time, at platform design and deployment. From then on, any application or service that utilizes the platform automatically benefits from the security policies already being put in place.

3.4 Leveraging Machine Learning and Artificial Intelligence

Manufacturing professionals have been optimizing operations and processes along the value chain for a long time. Traditionally, experts possessing intimate, proprietary knowledge of their particular machines, production cells, and assembly lines performed this optimization. In this people-centric analog environment, cross-domain orchestration and optimization is complex and typically does not scale to other production sites. With the connection of OT assets and the introduction of artificial intelligence (AI) and machine learning (ML) capabilities from the cloud, data can be used from one site to train ML models, and additional sites can take advantage of the resulting learnings. As new sites are connected and leveraged through algorithms and ML models, new data from those locations enhance the data set. Feedback communication is sent from the cloud-based models back down to the shop floor to implement changes based on these learnings. As ML models become trained and their results gain reliability, insights from the cloud trigger actions on the shop floor. Closed feedback loops generate virtuous cycles of learning, adapting, and efficiencies.

3.5 Consistent Device Management

A uniform approach to device management over the entire life cycle is required to ensure efficient administration of the many devices on the shop floor. The device lifecycle can be thought of as five discrete phases, each with its own requirements and goals.

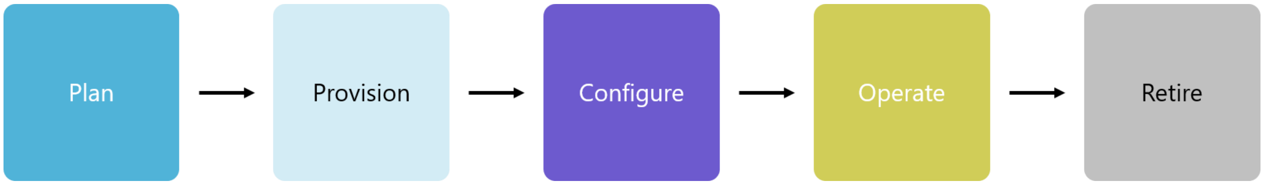

Figure 2: Typical Device Lifecycle

The planning phase includes developing an understanding of the expected use cases that drive the device capability requirements and is critical to both the individual device level and the system or solution level. This phase also includes elements of cost estimates, procurement, inventory, and implementation planning. Cataloging the devices in both brownfield and greenfield scenarios is an essential and very company-specific task with many dependencies needing consideration (e.g., network architecture and design, existing protocols, number of connected devices). Embracing open standards and automatic device discovery can be good solutions for speeding up this process.

Provisioning is the process of getting the hardware ready for deployment, including physical installation, integration into device management software, testing, and basic setup with other supporting systems such as an edge node.

Next, a device needs to be configured by setting different parameters to manage the software operations. This enables bulk updates, such as to the firmware, management of security, and ongoing operations.

- Download of software
- Setting the software parameter, e.g., parameters like endpoints, selection of data points, sampling, and publishing intervals
- Prepare interface on OT device, e.g., information model, data model, and message model
- Establish a connection to the device and cloud level

Operating the device includes collecting device telemetry, performing ongoing maintenance and management, and ensuring optimal performance. Holistic device and system monitoring are necessary to guarantee the greatest possible operational stability. In case of failure, it is essential to carry out an analysis with the help of a monitoring service. The flexibility to exchange hardware and software and update, reconfigure, and restore the connectivity solution with the help of a provisioning service is also fundamental.

The final stage is retirement, which is typically initiated by the end of the device service life, a device failure, or an upgrade cycle. Depending on the scenario, it will end up in a replacement or termination. In case of replacement, the faulty device should be replaced as soon as possible to avoid downtimes. Therefore, the device management service needs to configure the device automatically. In the event of device termination, information must be removed from the device management service, and firewall rules must be deactivated.

4 Types of Communication

The three different communication categories required to solve connectivity challenges are telemetry, control, and (device) management. Each one has unique properties, and the applicability of each varies according to individual solutions and level of communication within the network topography.

Each communication category can occur between the production assets & edge or edge & cloud (see Figure 3). We define this as the asset-edge interface or edge-cloud interface.

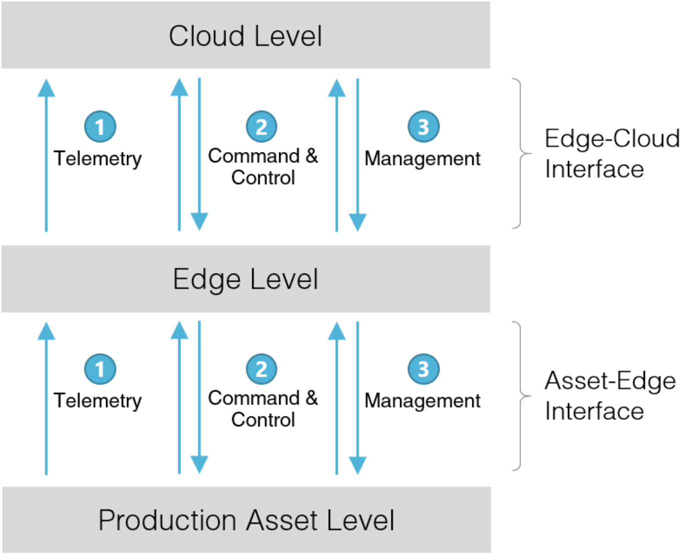

Figure 3: Three general communication types across three levels

The most common and prolific category is telemetry, which is the sending of raw data from the shop floor into the cloud. Examples include vibration readings in the range of milliseconds, pressure readings twice per minute, and device state status every 5 minutes. The messages typically contain data (e.g., current temperature) and metadata such as device ID and a timestamp.

Control data are product job information, recipe settings, or process parameters from central or regional servers or cloud hosts. Often control data originates and stays in the production asset level because of time criticality. However, hybrid edge and cloud architectures are being used to enable manufacturing use cases that were impossible prior to these new techniques, for example, automating control actions that previously required specially trained personnel.

Management data can be bidirectional and helps to control devices or machines. This typically includes metadata with the purpose of ensuring the entire system is operating correctly. Machine-to-machine communication is monitored, and information is utilized to ensure smooth operation. Examples include heartbeat monitoring as well as checking communication packets and their completeness.

5 Implementation

5.1 Production Asset Level

The primary purpose of the production assets on the shop floor is to run and optimize the manufacturing processes. They produce or are capable of producing massive amounts of data, and they are the primary source of telemetry data that is available within any level of the technology stack.

At the OT level, there are many different protocols. Specialized real-time protocols (like Profinet or Powerlink) are often not compatible with the IT systems since the first group often has different ISO / OSI Layers then the standard Ethernet protocol, commonly used by IT systems. For the connection between OT and IT levels, two standards have developed a wide adoption in the last years: OPC UA (Open Platform Communications Unified Architecture) and MQTT (Message Queuing Telemetry Transport). Both protocols are seen as a common factory standard since they are vendor and hardware independent. There are also open-source implementations for a variety of platforms. Most OT device manufacturers have included at least one of these two protocols in their state-of-the-art products. This is the reason why this working group mainly focuses on these two protocols.

OPC UA should be considered a connectivity implementation because of its widespread use in the manufacturing sector, self-describing behavior, data semantic standardization groups, and support for modern security standards. The OPC UA Foundation standardizes data exchange as a platform-independent, service-oriented architecture (SOA). Many forward-looking implementations such as OPC UA pub/sub and OPC UA over TSN are specified. However, this whitepaper focuses on the OPC UA client/server variant, as it is currently the most widespread. If the implementation of OPC UA is not reasonable or applicable, the use of MQTT as a transport protocol is a valid alternative. The choice of data formats and transport protocols should be considered based on the specific use case. For example, for a lightweight IoT sensor, a complex OPC UA implementation may not be possible due to hardware limitations.

Due to high adaptation costs in operational production facilities, there are still many legacy protocol standards, which have been in use for more than 20 years (e.g., Siemens S7 protocol). If OPC UA or MQTT is not natively supported, an adapter on the production asset level is a good solution to transform data from the proprietary asset interface to OPC UA or MQTT. An alternate solution is the installation of a software adapter at the edge level where the data is filtered, combined, and converted to the northbound edge protocol (MQTT). The overall goal is to transform and harmonize the data as early as possible in the communication process.

On the communication side, telemetry is typically passed to the higher network levels. Only minimal processing, filtering, compressing, and aggregating are done at the asset level. Instead, the machines are configured to efficiently pass information to other assets optimized for these compute-heavy functions, e.g., edge nodes.

In most circumstances, the shop floor machines are recipients of control and management information, acting on commands passed down from higher levels and generated from more complex and comprehensive models. Besides the “traditional” control mechanisms via PLCs, there is also the possibility of having a control model at the edge level. An example could be a machine learning model that evaluates telemetry data and sends control commands (e.g., updated parameters) via the edge adapter down to the asset. Device management communication is used to discover endpoints and specify which data points should be streamed in the telemetry messages. Also, the management of the device’s firmware lies in this domain. Security is the third aspect, with the possibility to set and change encryption mechanisms and security tokens (e.g., X.509 certificates).

Device management must keep the firmware of production assets up to date to secure the production assets, but the primary action necessary to minimize the attack vector is to use firewalls that deny access in general.

5.2 Edge Level

Edge nodes, whether installed on physical industrial PCs (IPCs) or a virtual edge device in the data center, enable a scalable, secure, and reliable architecture for connecting devices from the production asset level to the edge and cloud layer. Edge computing provides the benefit of applying logic at an immediate level. Different edge modules installed on an edge device take care of the functionality needed for a successful connectivity solution.

Communication modules ensure a secure and stable connection to the assets or other integration services. For modern protocols, like OPC UA and MQTT there are already various modules available, whereas proprietary protocols require vendor-specific communication modules. Additionally, cloud communication modules actively connected to the cloud are a vital feature of the edge level. Additional modules handle other functionalities that are necessary for an integrated solution: buffering, filtering, distribution of the messages to further endpoints, or other edge modules (e.g., AI modules), and preprocessing (e.g., message enrichment). The edge layer can run complex logic modules and act as a recipient of trained or updated machine learning models, applying the new logic to incoming data and making determinations from it. Device management modules for monitoring, updating, or security activities on the device are essential to managing large-scale edge ecosystems.

Due to the high volume of telemetry data, efficient processing and forwarding are essential. Filtering, compressing, and aggregating data at the edge can help to avoid overloading the network. Missing metadata is added, data formats get standardized, and short-term network failures can be bridged via buffer functionality. Due to the high volume of data, buffering or storing the data for an extended period on the edge device is not recommended. When making filtering and compression decisions, special attention should be paid to the potential future needs of applications and model training. Optimization for efficient transmission and storage based upon current data needs might result in missing or incomplete data sets for applications in the future.

Data might not be recoverable, and the future deficiency needs to be weighed against the current costs of transmission and storage.

Control messages can be received from any layer. A cloud service sends a new production job to the production asset or an AI model that analyzes the edge node's telemetry data, modifying the current production asset settings. Also, control messages from the production asset to an upper level, like alarm messages, are possible. Besides pure forwarding, control messages can also be intended for the edge device itself, e.g., config update. Due to the intermediate level, communication flows are vitally important, and the transmission of control messages needs to be reliable and secure. As a routing agent, the edge level needs to understand commands and route them throughout the system. If a control message needs to be acknowledged, control responses need to be mapped to the corresponding requests to route it back to the right sender.

Management data related to the production asset has to be interpreted and forwarded to the desired destination. Management data concerning the edge node must be routed between the device management edge modules and device management cloud services. A monitoring component provides logs and metrics, firmware and edge modules get updated, and certificates are renewed. Also, the settings need to be saved on a central service to restore them after a device outage or transfer it to a replacement device.

The edge level is inherently extending the network to remote site locations and increasing the number of possible attack vectors. Building a secure device is challenging, but challenges can be mitigated if we adhere to the principles and practices. Below are the attributes of secured and interconnected devices.

- Hardware-based root of trust: Physical measures resist side-channel attacks and protected by hardware
- Small trusted computing base: Private keys stored in the hardware protected vault and make it inaccessible to any software
- Compartmentalization: Hardware-enforced barriers between software components prevent a breach in one from propagating to others.
- Certificate-based authentication: Signed certificate, proven by unforgeable cryptographic key, proves the device identity and authenticity.
- Renewable security: Renewal brings the device forward to a secure state and revokes compromised assets for known vulnerabilities or security breaches
- Failure reporting: Any failure-probing security is reported to the failure analysis system

While edge devices typically provide the capability to run containerized software, the software management must be checking the base images for vulnerabilities and providing update paths that ensure smooth operations.

5.3 Cloud Level

Cloud services make use of flexibility, high computing power, and large storage capacities. In order to connect these services to shop floor equipment, a service is necessary that acts as the interface for receiving and sending messages to and from the lower levels. Thereby the data can be analyzed, visualized, and longtime stored. For example, cloud computing power can be used to train machine learning models based on data from the shop floor.

The Cloud Integration Hub is the cloud endpoint for any data receiving and sending to the lower levels, independent of the communication type. For security reasons, no upper level should actively establish a connection to a lower level. That’s why the Cloud Integration Hub is a passive service. Edge nodes or OT devices can connect to the service and get authorized by providing credentials. The Cloud Integration Hub doesn’t have to be one component. It might also be feasible that there are different components for different purposes and communication types.

For telemetry messages, a streaming-based solution optimized for high data throughput is recommended. Simple integration interfaces for other cloud services and an ability to distribute the data to different services should be supported.

As control data flows are bidirectional, two interfaces are necessary for sending and receiving control messages. The reliable transmission of the messages should be the focus.

Management data is received, sent, and stored by the Device Management Service. The Service provides provisioning, operating, and monitoring functionality for edge nodes and OT devices. It can install and update the firmware, manage the edge container, and supervise the security status of the devices. In addition, it stores all device-related information like device name, IP (Internet Protocol) address, configurations, installed edge container, and firmware status. If a device exchange is necessary, the information can be transferred to a new device.

From a security point of view, the cloud level provides centralized certificate management as well as identity and access control (IAC). Certificate management is essential to secure communication. The management needs to provide creation, renewal, revocation, and updated for all devices over the whole device lifecycle. The IAC is protected from unwanted modifications and can be used in combination with audit logs. As a best practice, the IAC should provide service principles to allow device management in unattended mode (e.g., to renew device certificates).

5.4 The interface between Production Asset and Edge Levels

The interface between production assets and edge level compute resources is usually a proprietary and vendor-specific implementation, dependent upon pre-existing installations and pre-programmed interactions. While some solutions are greenfield installations, this is not the predominant case for most manufacturers. Compute resources are generally limited on either side, and communication is optimized for efficiency and maintenance of the production line.

These factors mean any connectivity solution must be built within the context of the existing assets and limitations. Telemetry data is passed up to the edge level while command and managerial data is transferred back-and-forth to manage the operations and maintenance of the production level assets.

Due to the reasons given in a previous section, an OPC UA client/server is often the first choice for communication at this interface level. An edge device, via an OPC UA client interface, consumes the OT asset providing an OPC UA server interface. To pass telemetry data, the client can browse, read, and subscribe to OPC UA nodes and call OPC UA methods to obtain and transfer data. For commands, the edge device can also communicate with OT assets to perform functions such as browsing and writing OPC UA nodes and utilizing methods to make changes.

In many brownfield installations, both OPC UA and MQTT are not available, or it is not practical to upgrade the communication technology. In this case, adapters are needed to convert formats between proprietary protocols and promote standardization and interoperability. Adapters can be installed directly on the OT machines, and thus their presence and function are largely invisible to the rest of the system. This setup requires coordination between the asset vendor and the chosen target protocol. The adapter could also be installed on an IPC at an intermediate level between the production asset and edge layer, or it could be obtained through a marketplace and installed on the edge layer. Each setup requires different amounts of coordination with the asset manufacturer and forethought. Still, they all result in standardization and adherence to the OMP principle of interoperability beyond the edge level.

5.5 The interface between Edge and Cloud

Conveying data from the edge into the cloud via a streaming approach helps leverage independent cloud providers and drives a modular approach. Standardization takes place at the edge, and upstream computing can then be modularized to leverage best-in-class technology and capabilities. Plain MQTT or AMQP (Advanced Message Queuing Protocol) could be used, but a more forward-looking approach is OPC UA pub/sub. Standardization takes place at the edge, and upstream computing can then be modularized to leverage best-in-class technology and capabilities. OPC UA pub/sub, in contrast to plain MQTT or AMQP, does not require any transformation and therefore it is seen as the most promising candidate. Given the availability on this level, OPC UA pub/sub, which can use, for example, MQTT or AMQP, is the first choice for telemetry data because of the standardization, availability at the edge, and serialization capabilities.

With telemetry using OPC UA pub/sub, the edge provides an OPC UA publisher interface that can connect to one or more cloud-side MQTT/AMQP broker. In this case, clients connect to the MQTT/AMQP broker and consume messages. It is also possible to use a broker-less publisher-subscriber direct connection (using datagram protocols as UDP for transport).

If control from asset level to edge level is required, it is preferable that the edge level has an OPC UA server interface. If not possible, an approach where the asset sets values in its OPC UA server and the edge device (as an OPC UA Client) reacts on it is also feasible. The client needs to be notified as changes occur, e.g., by subscriptions or registered reads. For example, a data threshold for a machine learning model can be set on the server-side (asset). The client (edge) pulls these values and feeds the new threshold into its model; thus, the server indirectly controls the client by telemetry.

The predominant category of communication from cloud to edge is either command or managerial. Telemetry is not typically passed in this direction. The drivers of the data are the functions, logic, and intelligence applied in the cloud, which is then translated into command and management data to be passed down into the edge layer. The communication typically results in an action, either within the edge device or production assets. For example, parameters such as filters, logic, frequency, or timing could be altered based upon the learnings in the cloud.

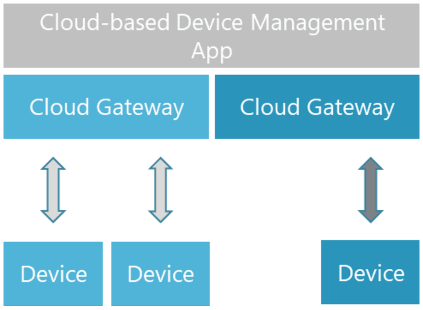

Figure 4: Cloud-based Device Management App

For example, as depicted in Figure 4, a device management solution provides management capabilities for edge devices. There are three edge devices from two different vendors that require specific communication. In order to abstract the device management solution away from the proprietary devices, the cloud gateways provide an open unified interface to control the edge devices. Examples of typical commands include start, stop, update firmware, renewal of certificates, download files as a trained machine learning model, and upload files.

The open interface receives commands from cloud applications while the propriety channel enables communication down to specific edge devices. The advantage of an open interface is that it is interoperable with any public cloud provider utilizing standard compute resources or cloud services. For the same reason, an edge container, enabling proprietary communication, should have an open interface to manage the connection settings.

6 Conclusion

From simple performance reporting to AI-optimized process control, IIoT connectivity is an essential cornerstone to digitalizing any industrial solution. For this, open standards and open-source implementations are important prerequisites.

The OMP establishes collaboration and open-source thoughts in the manufacturing domain, resulting in a robust approach that can scale easily throughout the technology stack, across manufacturing sites, and between different members within the value chain.

A complete technology solution includes connectivity as well as many other dimensions and technological considerations. To address this, the OMP has started different working groups covering various domains:

- Architectural guidance will be provided for standard manufacturing uses case through the Manufacturing Reference Architecture working group.
- Contextualization is imperative for transforming data into insight, and it is the main topic in the Semantic Data Structuring working group.

Subsequent publications by the IoT Connectivity Working Group will build upon this publication and dive more deeply into specific use cases and different layers of the technology stack. Topics such as device provisioning, certificate management, and store-and-forward mechanism for offline situations will be explored in-depth in future publications. The joint collaboration of different manufacturing companies will accelerate the development of production-ready software products and support the manufacturing community in implementing connectivity use cases that are reliable, innovative and secure.