OVERCOMING THE DATA DEADLOCK

Federated Learning in Industry

Challenges, experiences, and take-aways from a real-world implementation of federated learning in electronics manufacturing.

SIEMENS

KATULU

Contents

Introduction to Federated Learning

Federated learning is gaining momentum for industrial applications as it offers a practical solution for training machine learning models collaboratively while data remains secure on the local shopfloor. Hence, we can now optimize production across factories, organizations, and country borders.

Federated learning (FL), at its core, is a distributed machine learning (ML) paradigm that enables AI models to be trained collaboratively across a network of decentralized devices. Unlike traditional centralized models, FL keeps data local, thus preserving the integrity of proprietary designs, confidential production data, and intricate supply chain information. With data privacy paramount in today's regulatory environment, FL is a tool in the manufacturing industry's quest for AI-driven insights while safeguarding critical information.

This whitepaper presents an application of FL within electronics manufacturing, demonstrating enhancements in the quality control process of printed circuit boards (PCBs). The solution is used in two Siemens factories in Erlangen and Amberg with a ML model that detects anomalies in PCBs using process data. The FL model reduces the costs of manual quality inspection in both factories significantly.

How does federated learning work?

Every time you type on the virtual keyboard of your smartphone it collects data. At the same time, keyboard predictions are used to suggest words you might want to type next – and users experience improving word suggestions over time.

The personal notes, private messages, passwords, etc. you type are highly sensitive. This data is so sensitive that owning and storing it in a central storage would be significant business risk and challenging to explain to users. For this reason, FL comes into play to train a neural network for keyboard predictions without sending the data. This use case was the first commercial use of this technology.

Every one of us uses federated learning daily without knowing it.

This is how it works:

1. **Distribute initial AI model** – The initial model is distributed to the participating smartphones. This model is trained solely using internal datasets captured within the phone's testing system to predict mobile keyboards inputs. This first step is performed once per AI model.
2. **Collect and anonymize data** – While you type on your smartphone, training data is collected on the phone. Before the data is used for model training, it is anonymized by using privacy enhancing technologies such as differential privacy. Differential privacy adds noise to the data without changing the meaning of the data, like putting frosted glass in front of the data. Applying data anonymization is optional but recommended for most cases.
3. **Train locally** – While you are using the keyboard, training data is collected only on the phone. The AI model is trained asynchronously at night when the device is not used and plugged in to a charger. This training happens directly on the device, ensuring that personal data remains private and never leaves the phone.
4. **Create model update locally** – After the local training, each device generates a model update. Instead of raw data, this update contains summarized changes to model parameters (weights) based on your typing habits. It is a compact and anonymized representation of improvements.
5. **Transmit local update** – This compact representation of model changes is sent to a central server when the device is idle, plugged in, and connected to Wi-Fi.
6. **Learn globally** – A cloud-based FL backbone combines all device model changes to a new global model. This can use different algorithmic strategies (e.g., federated averaging).
7. **Deploy updated global model** – Every participating smartphone receives an updated global model, so every user enjoys enhanced and more accurate keyboard predictions.

In the outlined example, all participants are performing the identical task of "mobile keyboard predictions" using the same input data type of "typed text." This type of federated learning is called horizontal federated learning (HFL), where multiple parties have the same features for different products.

FL is an established technology used in smartphones. It can be used in many more ways in industry – from discreet manufacturing, process industries, mobility to health care – wherever AI can optimize systems across factories, organizations, and country borders without having to share private data.

How does federated learning help in industry?

The transition of artificial intelligence (AI) from research to industrial application is a challenge that often fails for many reasons – ranging from regulatory barriers to strategic concerns, technical issues, and cybersecurity risks. FL is a compelling option to addressing these challenges.

Regulatory barriers – Export control includes production data of dual-use goods (military and non-military). Financial insider information for publicly listed corporations can be regulated, think of timestamps providing insights into production through-put, therefore company performance. And occasionally even General Data Privacy Regulation kicks in as a challenge in industry.

FL keeps the data locally and processes it at the source thus ensures compliance while leveraging any data across country borders, organizations, and factories.

Strategic concerns – Real-world use cases often require much more data than a single organization or data owner can provide. Because different data owners have different interests and concerns, sharing data can be a sensitive issue. Even in collaborations without direct competition, concerns about unique selling proposition and intellectual property prevent data sharing. Geopolitical tensions intensify these concerns.

A FL setup reflects the strategic role of data. While data remains with its original owner and is processed just locally, collaboration partners agree to share just specific insights among them, represented by the AI model. Consequently, businesses can harness the power of industrial AI without compromising their trade secrets or intellectual property.

Cost concerns – Modern manufacturing sites generate enormous amounts of data, especially when machine vision or high frequency data is involved. While data transfer cost to the cloud may seem minimal at the first glance, the sheer volume drives costs up – and: data cloud transfers add carbon footprint.

FL reduces data transfer costs and the associated carbon footprint.

Cybersecurity risks – Central pools of combined data are an attractive target for malicious actors, increasing the damage caused by data breaches. A party operating a central data pool must be able to deal with cybersecurity risks and potential legal and reputational ramifications.

With FL, a central data pool is not required anymore, an elegant way to mitigate associated risks.

As outlined, federated learning is emerging as a compelling solution addressing the challenges of industrial AI and promises a new privacy-preserving and collaborative era for industrial AI integration.

How can federated learning work in a global value chain?

Let us imagine two companies: Browsing, Inc. has data on browsing habits, and Shopping Corp. possesses data about shopping habits. Their customer base overlaps.
While both companies have data on the same customer (or the same set of entities), the data attributes (“features”) they possess are different. With vertical federated learning (VFL), they can collaboratively train a model using their combined features without directly exchanging raw data about their customers.
For VFL, the model is basically cut into pieces and each company trains their piece or model layers, including a “cut layer,” that represents the connecting interface between the model parts.

VFL can be applied to whole supply or value chains, where every participant trains a model part for their part of the chain. Applying this technology to industrial settings could yield significant improvements. In production or quality inspection processes, for example, the trained models can be used to detect anomalies in the products or predict quality of work pieces and the knowledge for detecting certain types of anomalies is gained from other parties.

What is needed to implement federated learning?

With great excitement for its potential, FL is a deep-tech combination of mastering ML and mastering distributed systems. But what are the requirements for deploying industrial-grade FL?

1. **Privacy and data security** – Protecting the data diligently requires the prevention of data reverse-engineering from the AI model. This is where differential privacy helps as it adds sufficient amounts of 'noise' to camouflage real data. However, the difficulty lies in calculating and applying the right amount of noise as it must not change the meaning of the data, and it must be adapted to the number of operations performed on the model.

2. **Heterogenous data** – A major difference between the smartphone FL and industrial use cases is that smartphone data is homogeneous and the smartphone itself a closed system. In contrast, every factory setup is different and industrial data from each entity can vary in various aspects, e.g., available data points, formats, intervals, types, and even the storage where the data is persisted.
The capability to train the same model with different data is essential for a successful real-world industrial implementation.

3. **Combining machines, processes, and materials** – When developing ML models for industrial applications, the possible combinations of machines, processes, and materials that can be covered with a single model must be considered.

However, it is often unknown beforehand which combinations can be grouped into one model together. In the case of FL, detailed information about machine setups and processes are kept confidential to maintain the privacy promise, which makes it challenging to define cohorts of comparable combinations. As a result, models cannot deliver optimal performance and the power of transfer learning remains untapped.

4. **Device and setup heterogeneity** – The heterogeneity of involved devices and setups presents a significant challenge in industrial settings as hardware specifications vary as well as data collection methodologies and operating conditions. Aggregating and harmonizing data from these devices, as well as designing methods that can effectively generalize across the different setups.
In addition, local FL software components ("FL Agents") must provide a high adaptability to their infrastructure environment.

5. **Initial model** – The discussed scenarios require an initial model to start off with, which implies that lab data is available for development of the first model. However, there are cases in which organizations want to provide digital services based on FL and lack sufficient data to create an initial model. This is particularly true for vertical FL cases where every party performs a different task. But also, horizontal cases with a high level of combinations can require data from other parties to develop the first model, which is a blocker for most FL frameworks.

6. **Data and model monitoring** – In a FL setup, two different types of monitoring are required: local monitoring for inference and privacy-preserving federated monitoring for model training. At the facility of each participant, a local component monitors local model drifts and biases.
The challenge however comes with the privacy-preserving federated monitoring for model training because certain model drifts, biases, data gaps and poisoned data sets cannot be detected, if models and data sets of each participant are analyzed isolated at each site.

7. **Decentralized compute power** – Locally performed model training requires decentralized AI compute resources. So far, all small footprint AI hardware accelerators support inference only (e.g., Intel's Myriad, Google Coral, etc.), not training.

Industrial applications for federated learning

FL establishes trustworthy learning systems across several parties. The following application areas illustrate some of the potential that FL can unlock today.

1. **Predictive maintenance** – Collectively analyzing equipment performance data across manufacturing sites helps minimize downtime and optimize maintenance schedules. Collaboration can be realized between partners with similar processes or production assets in heterogeneous processes.

Similar equipment manufacturers (OEMs) can deploy FL agents (e.g., analytic apps) to their customers to gain knowledge from data they usually cannot access because of confidentiality (such as CNC machines used in defense manufacturing). Resulting models trained with FL can be used for product improvements or value-added services (e.g., a detected machine anomaly triggers a field service action by the OEM).

2. **Quality control and defect detection** – Automatically detecting defects requires a vast amount of high-quality data to teach an AI model the nuances of every defect in any given situation. FL empowers quality control by leveraging data from multiple plants, overcoming the lack of data to enhance defect detection models. This ensures that quality is assessed automatically, and defects are found without manual inspection. With federated quality control, strict quality standards can be achieved, and costs of quality can be effectively reduced. In the case of optical inspection, FL significantly decreases data transfer costs because the vast amount of visual data generated across different plants is processed and kept on the edge.

3. **Predictive quality control** – Instead of evaluating quality at the end of the process, predictive quality control goes a step further by leveraging data points from the entire process to predict quality at the earliest stage possible. With the right FL platform, data from the entire process can be used for model training. For the inference, only features that are available early in the process would be used to predict quality at earliest stage.

This avoids unnecessary process steps for faulty parts and allows optimizing the process accordingly. As a result, a higher percentage of products meets quality standards on their first attempt, leading to an increase in first pass yield and reduced manual rework.

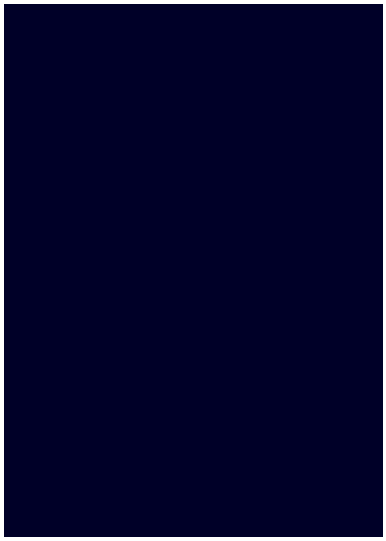

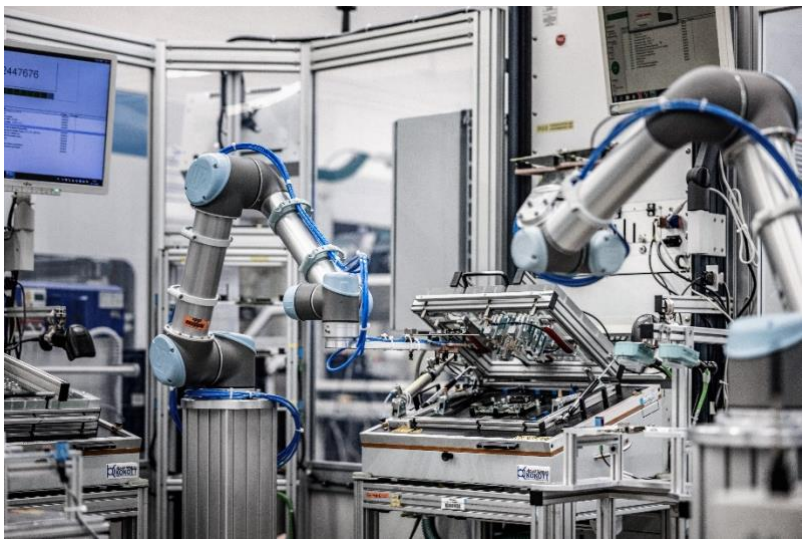

4. **Process optimization with internal best practices** – FL enables the analysis of production processes across different plants to identify best practices and optimize efficiency while maintaining confidentiality. Internal benchmarking of each plant's performance with cross-border data transmission safeguards the intellectual property of each location. Once effective practices are identified, they can be deployed globally while protecting trade secrets of each location and complying with regulations, including export control. This approach facilitates a comprehensive understanding of process strengths and areas for improvement, guiding the identification and implementation of worldwide best practices while maintaining data integrity and privacy.

5. **Process optimization for value chains** – Optimizing a process across organizations is challenging and requires precisely defined interfaces. With the use of VFL, process parameter optimization can advance in process planning throughout the value chain. By

training models on partners' data, insights into their process steps are obtained without direct access to data. Integrating insights from partners' process steps achieves more effective planning outcomes. This holistic approach overcomes the limitations of isolated step planning, ensuring every interdependency is acknowledged and optimized. Its outcome is a more efficient and cohesive process that minimizes inefficiencies and enhances collaboration among partners. This approach to process optimization could elevate the entire value chain, leading to superior performance and outcomes.

6. **Collaborative research** – FL allows researchers from different companies to join forces and train ML models on data that they cannot see. For example, researchers from different pharmaceutical companies use FL to train a model to predict the effectiveness of new drugs. Partners within industrial value chains can research ways to improve the sustainability of the entire chain by taking a holistic approach. Each company keeps their own data private, but they could share the model parameters to train a more accurate model.

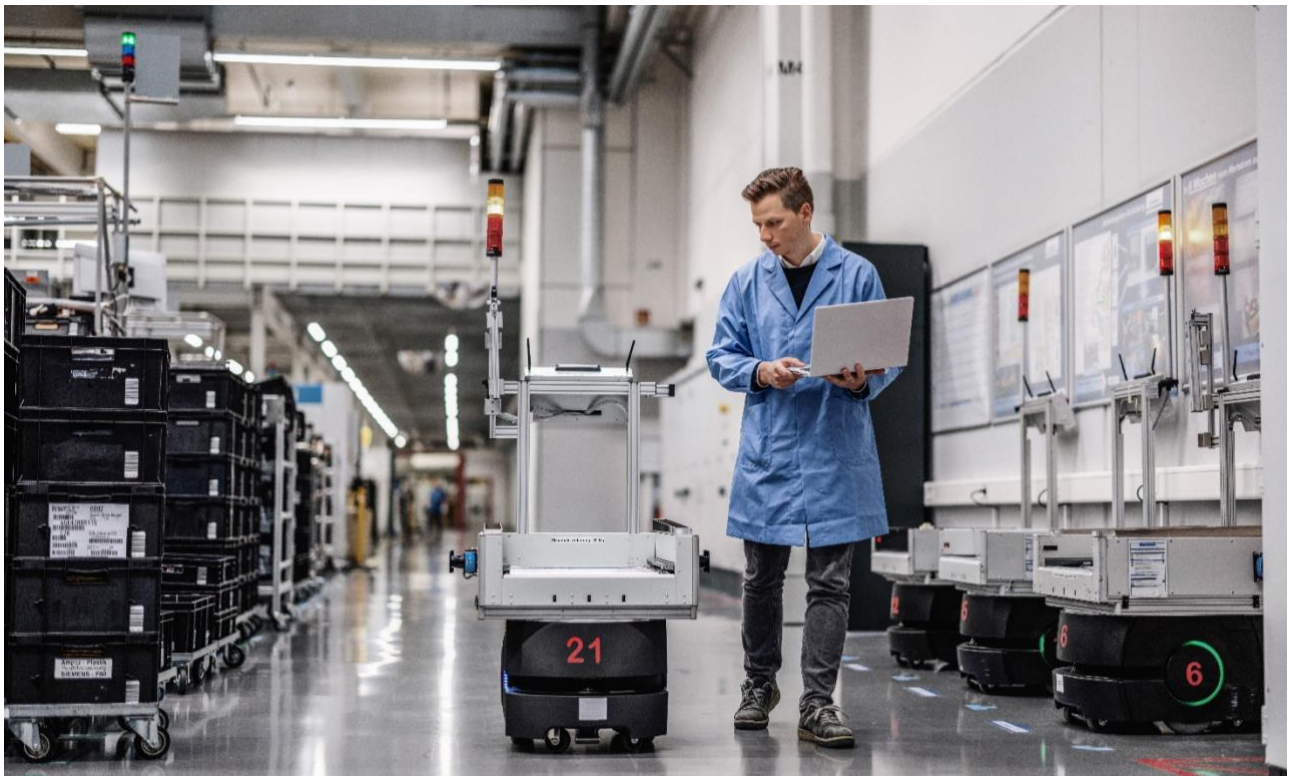

The PCB Production Process

Siemens factories in Erlangen and Amberg produce SIMATIC PLCs and SINAMICS drives for factory automation and motion control. The PCBs used in these products are built using the following production process:

1. Solder paste is printed on the PCB through a product-specific stencil.
2. The amount and position of transferred solder paste is inspected for each individual component contact.
3. The electrical components are placed on the board (between 50 and 1,500 components per board).
4. Reflow soldering of the PCB with product-specific temperature profiles is performed.
5. The automated optical inspection (AOI) identifies incorrect component placement, solder bridges, or open solder joints.

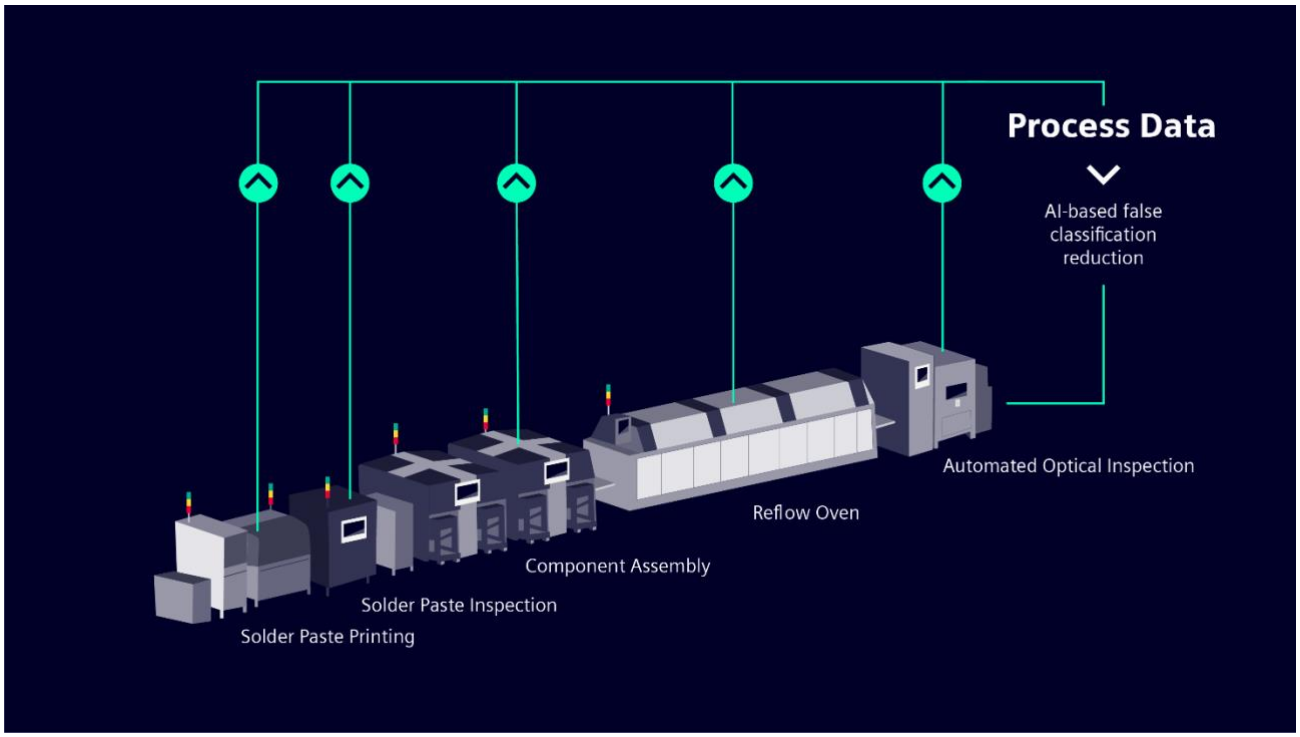

Figure 1: A typical PCB production line with automated optical inspection (AOI)

The PCB assembly process is measured by how well it produces good boards on the first try. Quality problems should be identified quickly to guarantee a high output from the line, so a proper classification of the AOI is essential for quality.

Our Implementation

At the Siemens factories, the AOI performance was insufficient and led to error slippage and false calls of good PCBs. Knowing the technological limitations and striving for a scalable solution, our Siemens factory digitalization team had been using an XGBoost-based AI-model to reduce false AOI classifications by using the process data along the whole production line for quite some time. However, XGBoost sticks to the PCB variants it has been trained on, thus does not generalize, and needs adoption to each different PCB variant, production line, and location.
The collaboration project of Siemens and Katulu introduced FL to further increase the performance of the AOI and reduce false calls, laying the foundation for a rollout to additional Siemens factories, including those in China.

Trust requires a trustworthy stack.

The IE provides many building blocks to consistently govern and manage heterogenous IT workloads in OT environments as depicted in Figure 2:

- Data bus to quickly establish data connections into the OT world.
- AI Software Development Kit to prepare and package self-contained AI pipelines including trained AI models, custom configurations, as well as package dependencies to make them ready for rollout.
- AI Model Manager to keep track of AI model versions deployed and to orchestrate their execution on the shop floor.
- AI Inference Server to reliably execute AI models on IE devices and wire them to data-sources as needed.
- AI Model Monitor to acquire system and model metrics during operations on the shop floor.

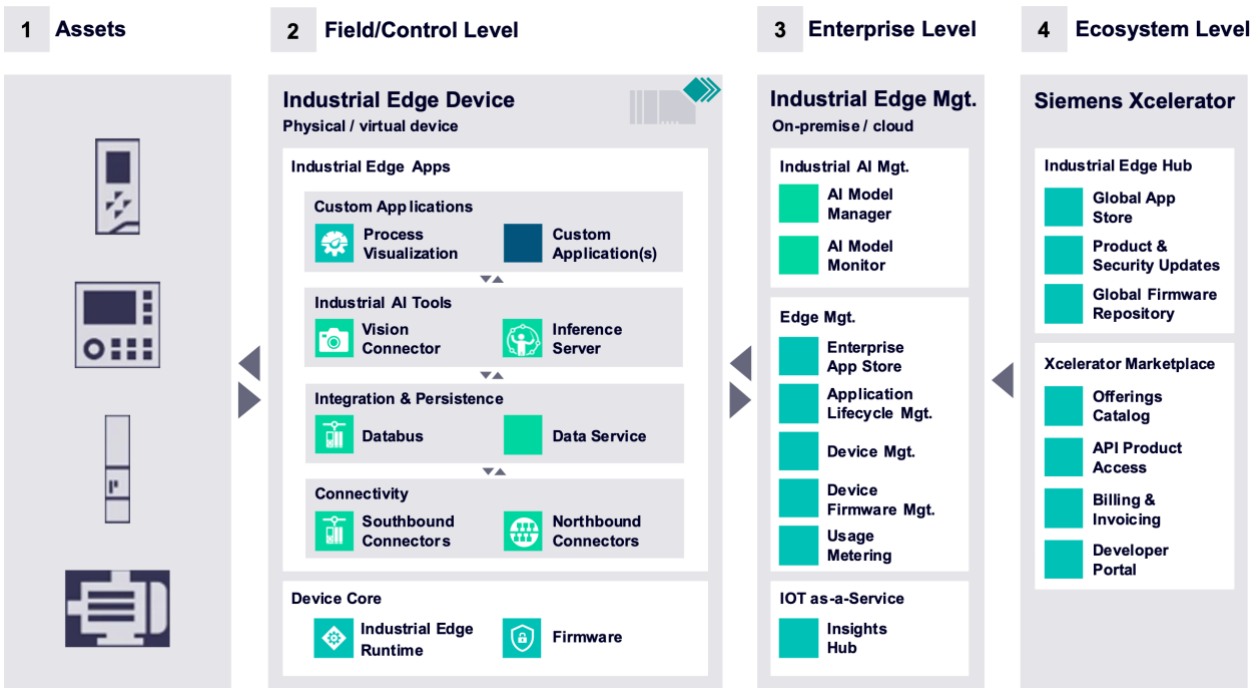

Figure 2: Siemens Industrial Edge building blocks

The main layers of our implementation for our factories (see Figure 4) were

- the *field level* with data collection capabilities
- the *factory level* with training capabilities from the Katulu Agent
- the *private cloud* with deployment capabilities
- the *public cloud* running the Katulu Platform with an FL server for model aggregation and sharing capabilities.

Each layer comprises software components that can be installed on each site and on every client using Siemens Industrial Edge downloadable tools.
The Industrial Edge provides connectors to integrate data from diverse data sources (e.g., OPC UA). Using the Katulu Agent, our project added a FL client application to the IE platform which serves to harmonize the diversity among various sites, the involved hardware, and their respective data infrastructures.

The *field level* consists of at least one edge device that receives data from the underlying (manufacturing) process (e.g., via OPC servers collecting variables from PLCs).

The Industrial Edge app runs on top of the Siemens Industrial OS, which is a Linux distribution hardened for industrial purposes. The Edge Runtime is best seen as an orchestrator that runs and monitors container-based software components, including connectors for Southbound Connectivity (i.e., OPC). The Data Collection & Aggregation module helps to preprocess data coming in from different sources, e.g., via the IE DataBus. The IE DataBus also serves to stream preprocessed data northbound.
The AI inference server continuously applies models on ingested data as they are deployed from higher levels.

The *factory level* hosts an IE device with sufficient resources with respect to memory, storage, CPU and GPU to train local models using the Katulu Agent. We have paired the factory model training component with a factory-level database, fed by the field level with preprocessed data.

Collaboration with other Katulu Agents is realized by connecting to the Katulu FL Server as part of the Katulu Platform in the *public cloud*. To secure the traffic, a local firewall is deployed, and all means of transport encryption used – reducing leaking of model information during the FL process.

To register multiple models for the factory (e.g., as a result from downloading the global FL model from the public cloud), the AI Model Manager is used. This component serves all edge devices on the field level with models suitable for local inference. By selecting the respective model, a one-click deployment is triggered. On the private cloud layer, we provide a CI/CD pipeline for bridging the gap between the global model and the factory. The pipeline downloads new models respectively model versions from the Katulu platform. A validation step is included ensuring sufficient model quality on local factory datasets. Subsequently the model is packaged to provide a deployable artifact that can be registered at the AI Model Manager and can finally be executed on the AI Inference Server in the field.

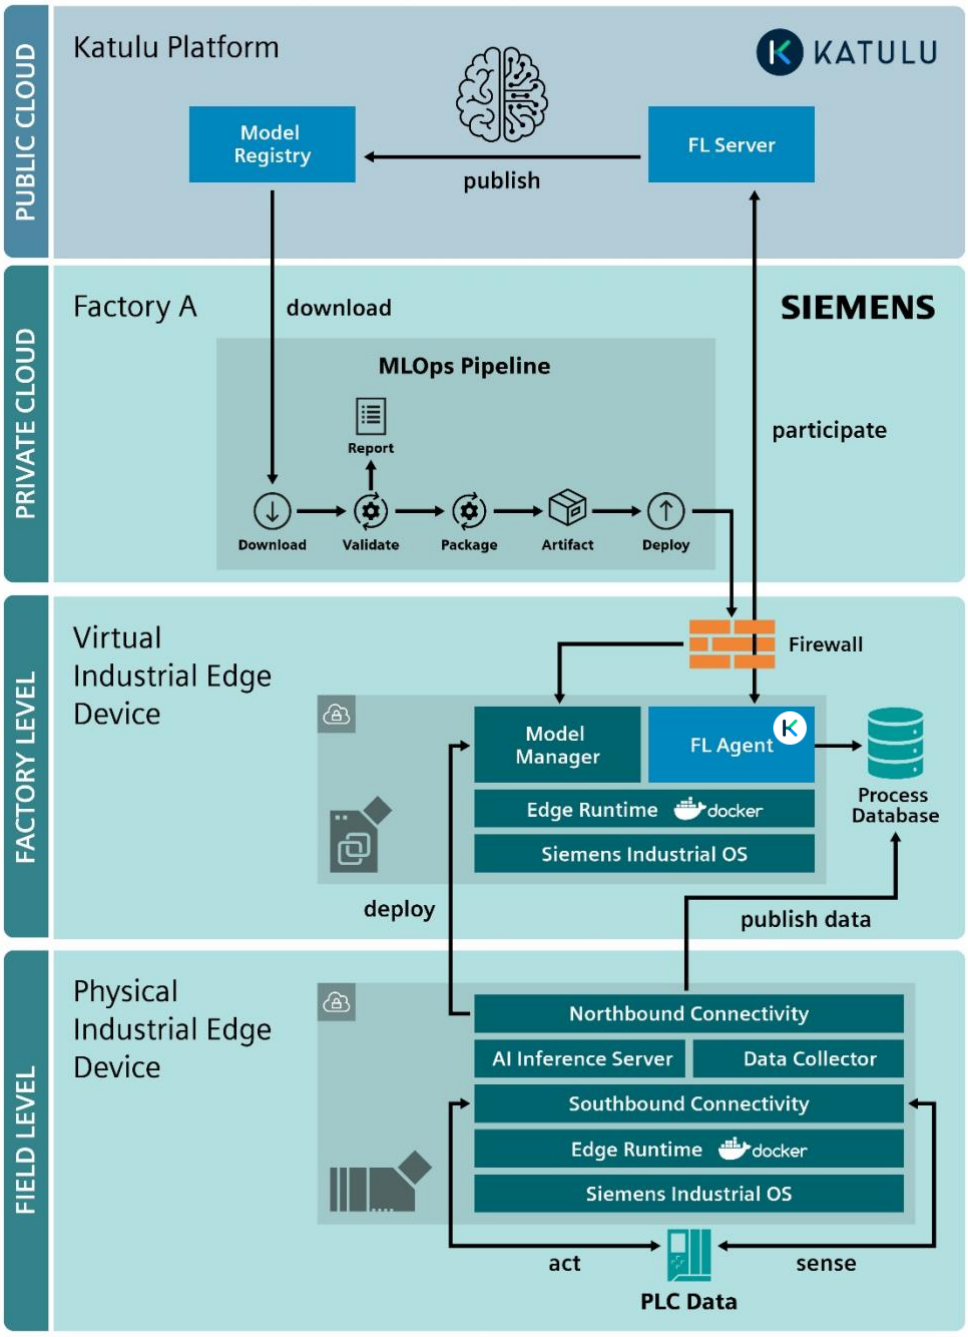

Figure 3: Architecture and technology stack for one client (factory) to participate in FL

The Katulu Platform addresses the challenge of privacy and security in FL systems. A prominent method is differential privacy, which ensures that client data is not disclosed in any form. In particular, the Katulu Agent in the factory adds noise to the uploaded model weights, which prevents potential reverse engineering to actual model parameters and especially raw data that is used for training the local model.

The integrity of the communication between the FL server and all participating FL Agents is ensured with M2M OAuth 2.0. This mechanism addresses the identified security challenge and enforces that all clients need to authenticate at the Katulu platform with their pre-signed tokens, they have received initially.

In Figure 4 we depict the overall interplay between two factories Gerätewerk Erlangen and Elektronikwerk Amberg.

Both factories implemented the complete architecture on every level and are linked by the Katulu Platform for a joint FL model.

The approach can be easily scaled to further factories using Siemens IE.

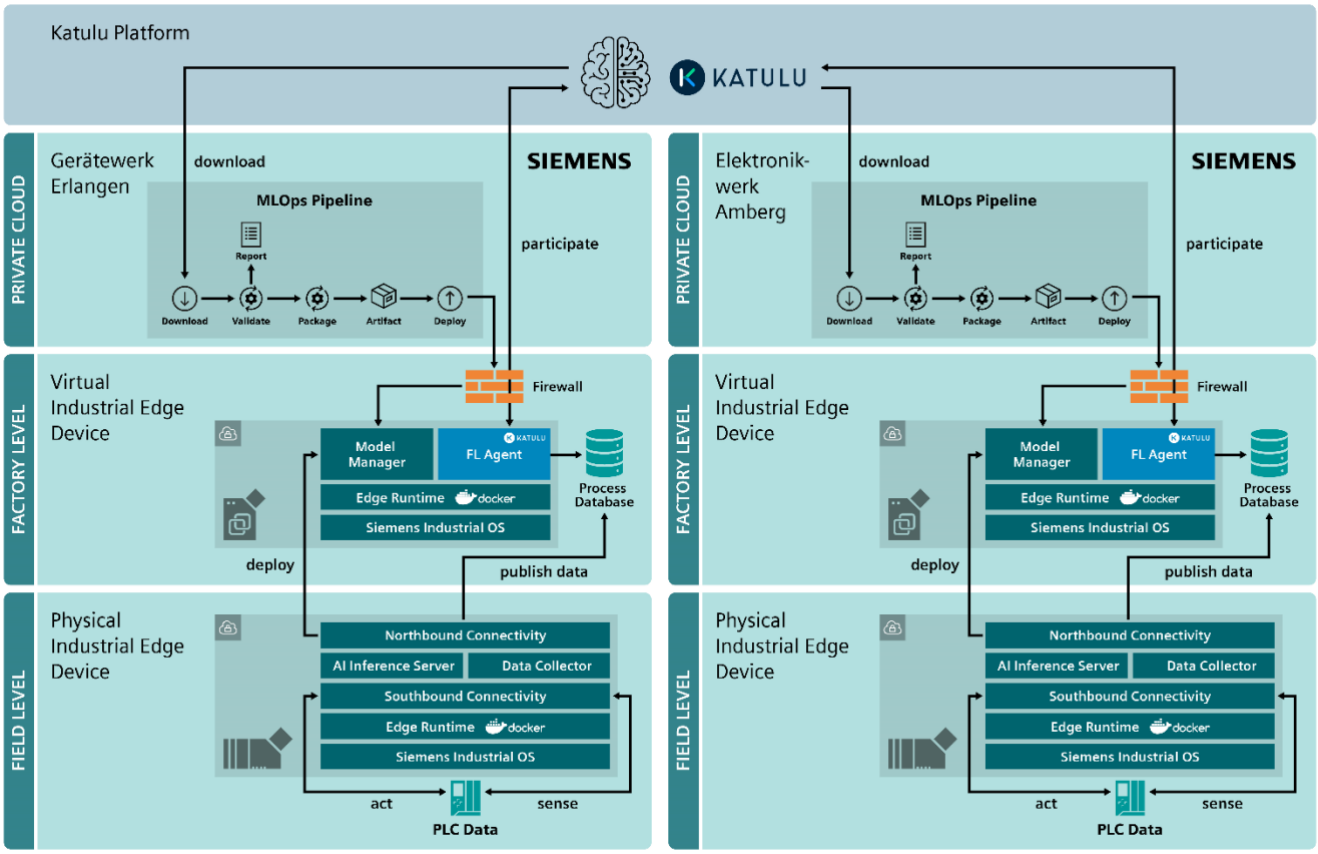

Figure 4: FL architecture for training models across two factories

Tackling the challenges

Let us look at the FL implementation challenges and how we address them jointly with our trusted stack:

1. **Privacy and Data Security** – Manufacturing data is defined as production technology and therefore intellectual property and export control relevant. As some of our products can be used for critical applications, we must control access to this data. Typically, we apply the need-to-know principle and try to keep the data locally within the factory intranet. This also helps to protect the data against manipulation or unwanted deletion, but when the data is used to train an AI model other mechanisms like Differential Privacy are needed. Katulu automated differential privacy and took out the complexity of this approach. The automation of this privacy enhancing technology ensures that the right amount of noise is always applied, and the amount of model operations is limited to prevent model attacks. Privacy of sensitive industrial data requires only the right tools, which made the usage also very simple for this implementation.

2. **Heterogenous Data** – The different productions lines differentiate in terms of manufacturer, model type and generation, which leads to different input data. Furthermore, all Siemens factories oversee their own IT infrastructure, which leads to different data storages including schemas. The capability to train models in such environments is essential to succeed with real-world industrial AI use cases. Katulus platform always ships the model together with the needed data transformation pipeline, which is why the resulted heterogenous input data can be transformed by the Katulu Agent. This approach offers the flexibility needed for iteratively federated model development, which decreases time-to-market and integration costs. The Katulu pipeline can also be used together with ETL pipelines already implemented with Siemens industrial Edge environment (e.g., using Siemens Mendix) as is the case with the Siemens factories.

3. **Combinations of Machines, Processes & Materials** – The high number of different PCB variants regarding layout, type, and number of components cohort management can lead to AI models without optimal performance, because the different groups of PCBs have different optimization points. This can be addressed with Katulus Cohort Management, which identifies automatically which PCBs can be trained together in the same cohort. The model of the factories in Erlangen and Amberg did not yet require a dynamic grouping in cohorts to improve model performance. However, the global rollout will expose the model to more combinations with different attributes, which is why we have the appliance of this functionality planned.

4. **Device and Setup Heterogeneity** – The Siemens factories – not uncommon for industrial environments – use a diverse range of manufacturing equipment, each with varying hardware specifications, data collection methodologies, and operating conditions. However, Siemens IE provides a wide range of connectors to integrate data from diverse sources and industrial protocols (e.g., OPC UA). The Katulu Agent has the streaming and query capabilities to leverage the connected data source. The Katulu Agent, which was added to the IE platform as an FL client application, served as a harmonization layer across various sites, hardware, and data infrastructures.

5. **No initial Model** – The Katulu platform provides AutoFL for horizontal and vertical federated learning use cases to develop initial models without any access to the data. Katulu developed its unique approach to neural architecture search for federated learning to find the best performing model on distributed private data. Since the initial data sets of the two factories were provided to Katulu in advance, an initial model was developed with access to the data. However, the described AutoFL solution of Katulu can be combined with the Siemens IE environment.

6. **Data and Model Monitoring** – For the time of the implementation of this project we needed to use custom tooling for the data and model monitoring. But Siemens will release their model monitoring suited for model inference in Q1 2024. Katulu will cover the data and model monitoring for the federated model training. This privacy-preserving federated monitoring for model training will be able to detect sensor and model drifts, biases, data gaps and poisoned data sets, which cannot be detected, if models and data sets of each participant are analyzed isolated at each site.

7. **Absence of User-Friendly FL Collaboration Tools for Industrial Clients** – To address the challenge of providing a user-friendly FL approach on a decentralized infrastructure of potentially multiple sites and industries, the different components like the Katulu Agent are provided in the Siemens edge ecosystem's app store and can be installed with a one-click deployment. After successfully installing the Katulu agents and the respective connectors, the data collection process can be started, and the alignment of the common FL data scheme and the global initial model can be initiated between the involved parties.

8. **Decentral compute power required** – The Siemens hardware for the Industrial Edge environment ranges from small edge devices up to GPU-powered industrial edges and the flexibility to virtualize an edge on a server. Since Katulu supports IE, the hardware for the Katulu Agent can always be sized according to requirements and easily adjusted in requirements change. Since the initial hardware requirements of the use case were unclear, we started with a virtual edge on a server. For the global rollout the virtual edge can be replaced with a dedicated hardware-based IE.

Our implementation results

After fine-tuning the AI model, FL clearly outperformed the central data approach (see Table 1). We see a better generalization capability and increased robustness against local noise in the data. This is why the initial F1-Score of 84% increased quickly in operations to 98% and Recall from 77% to 95%.

We are now able to scale our solution faster to other factories and, in addition, avoid orchestration and normalization of centralized training data

Table 1: Model Performance: At start, a central data and model would have outperformed FL, however the excellent FL generalization let very soon to outstanding model performance.

	local data & local model		**central data & central model**	**local data & federated model**
	Fact. GWE	Fact. EWA		
F1-Score	84% (+0%)	76% (- 8%)	88% (+4%)	**84%** (↗ 98%)
Recall	73% (- 4%)	66% (- 11%)	82% (+ 5%)	**77%** (↗ 95%)

We have laid the foundation for a further rollout of the solution – even into geographies with traditionally high hurdles for data access. Siemens is proactively looking for electronic manufacturers interested in joining this use case.

There are several model architectures (e.g., XGBoost) that are not compatible with federated learning. As our factories used XGBoost before, we had to train a new model (DNN).

For our specific tabular data-based use case, the decentral need for additional compute resources was limited. We also understand, the needs for decentral compute power are difficult to predict, hence need experimentation.

Katulu and Siemens see this collaboration as a blueprint for successful ecosystem partnerships.

We have integrated the Katulu FL platform smoothly with the Siemens Industrial Edge environment, combining two best practices for trusted analytics to an industrial-grade, full-stack solution.

Where can I get federated learning?

Katulu GmbH, a startup for industrial AI from Hamburg, is an innovation leader for industrial federated learning. As a pioneer in this field, the Katulu team has many years of cutting-edge experience and developed a patented approach to make federated machine learning applicable for real-world scenarios, where data and machinery are heterogeneous. Katulu technology enables leading companies to learn across their factories, machines, customers or along their value chains. Katulu offers a software solution for federated learning specifically for industrial applications and complementary services like consulting, development, and operations. With these capabilities Katulu has become the trusted partner of major players within the industrial sector. Further information is available on the Internet at www.katulu.com.

Siemens Digital Industries (DI) is an innovation leader in automation and digitalization. Closely collaborating with partners and customers, DI drives the digital transformation in the process and discrete industries. With its Digital Enterprise portfolio, DI provides companies of all sizes with an end-to-end set of products, solutions, and services to integrate and digitalize the entire value chain. Optimized for the specific needs of each industry, DI's unique portfolio supports customers to achieve greater productivity and flexibility. DI is constantly adding innovations to its portfolio to integrate cutting-edge future technologies. Siemens Digital Industries has around 72,000 employees internationally.

Siemens AG (Berlin and Munich) is a technology company focused on industry, infrastructure, transport, and healthcare. From more resource-efficient factories, resilient supply chains, and smarter buildings and grids, to cleaner and more comfortable transportation as well as advanced healthcare, the company creates technology with purpose adding real value for customers. By combining the real and the digital worlds, Siemens empowers its customers to transform their industries and markets, helping them to transform the everyday for billions of people. As of September 30, 2022, the company had around 311,000 employees worldwide. Further information is available on the Internet at www.siemens.com.

Partnership

This whitepaper has been written by Thomas Blumauer-Hiessl, Boris Scharinger, Thomas Kaufmann, Erik Schwulera, Konstantin Schmidt, Michael Kuehne-Schlinkert who genuinely enjoyed working together.

Contact us to learn more:

Michael Kuehne-Schlinkert

michael@katulu.io

KATULU

Anne Mareike Schlinkert

anne@katulu.io

KATULU

Erik Schwulera

erik.schwulera@siemens.com

SIEMENS

Boris Scharinger

boris.scharinger@siemens.com

SIEMENS

Veröffentlicht von
Katulu GmbH
An der Alster 6
20099 Hamburg
Germany
Um weitere Informationen zu erhalten, wenden Sie sich bitte per E-Mail an:
michael@katulu.io