International Journal of
Computing and
Engineering
(IJCE)

Federated Learning for Ransomware-Resilient Industrial IoT: A Decentralized Framework for Secure AI at the Manufacturing Edge

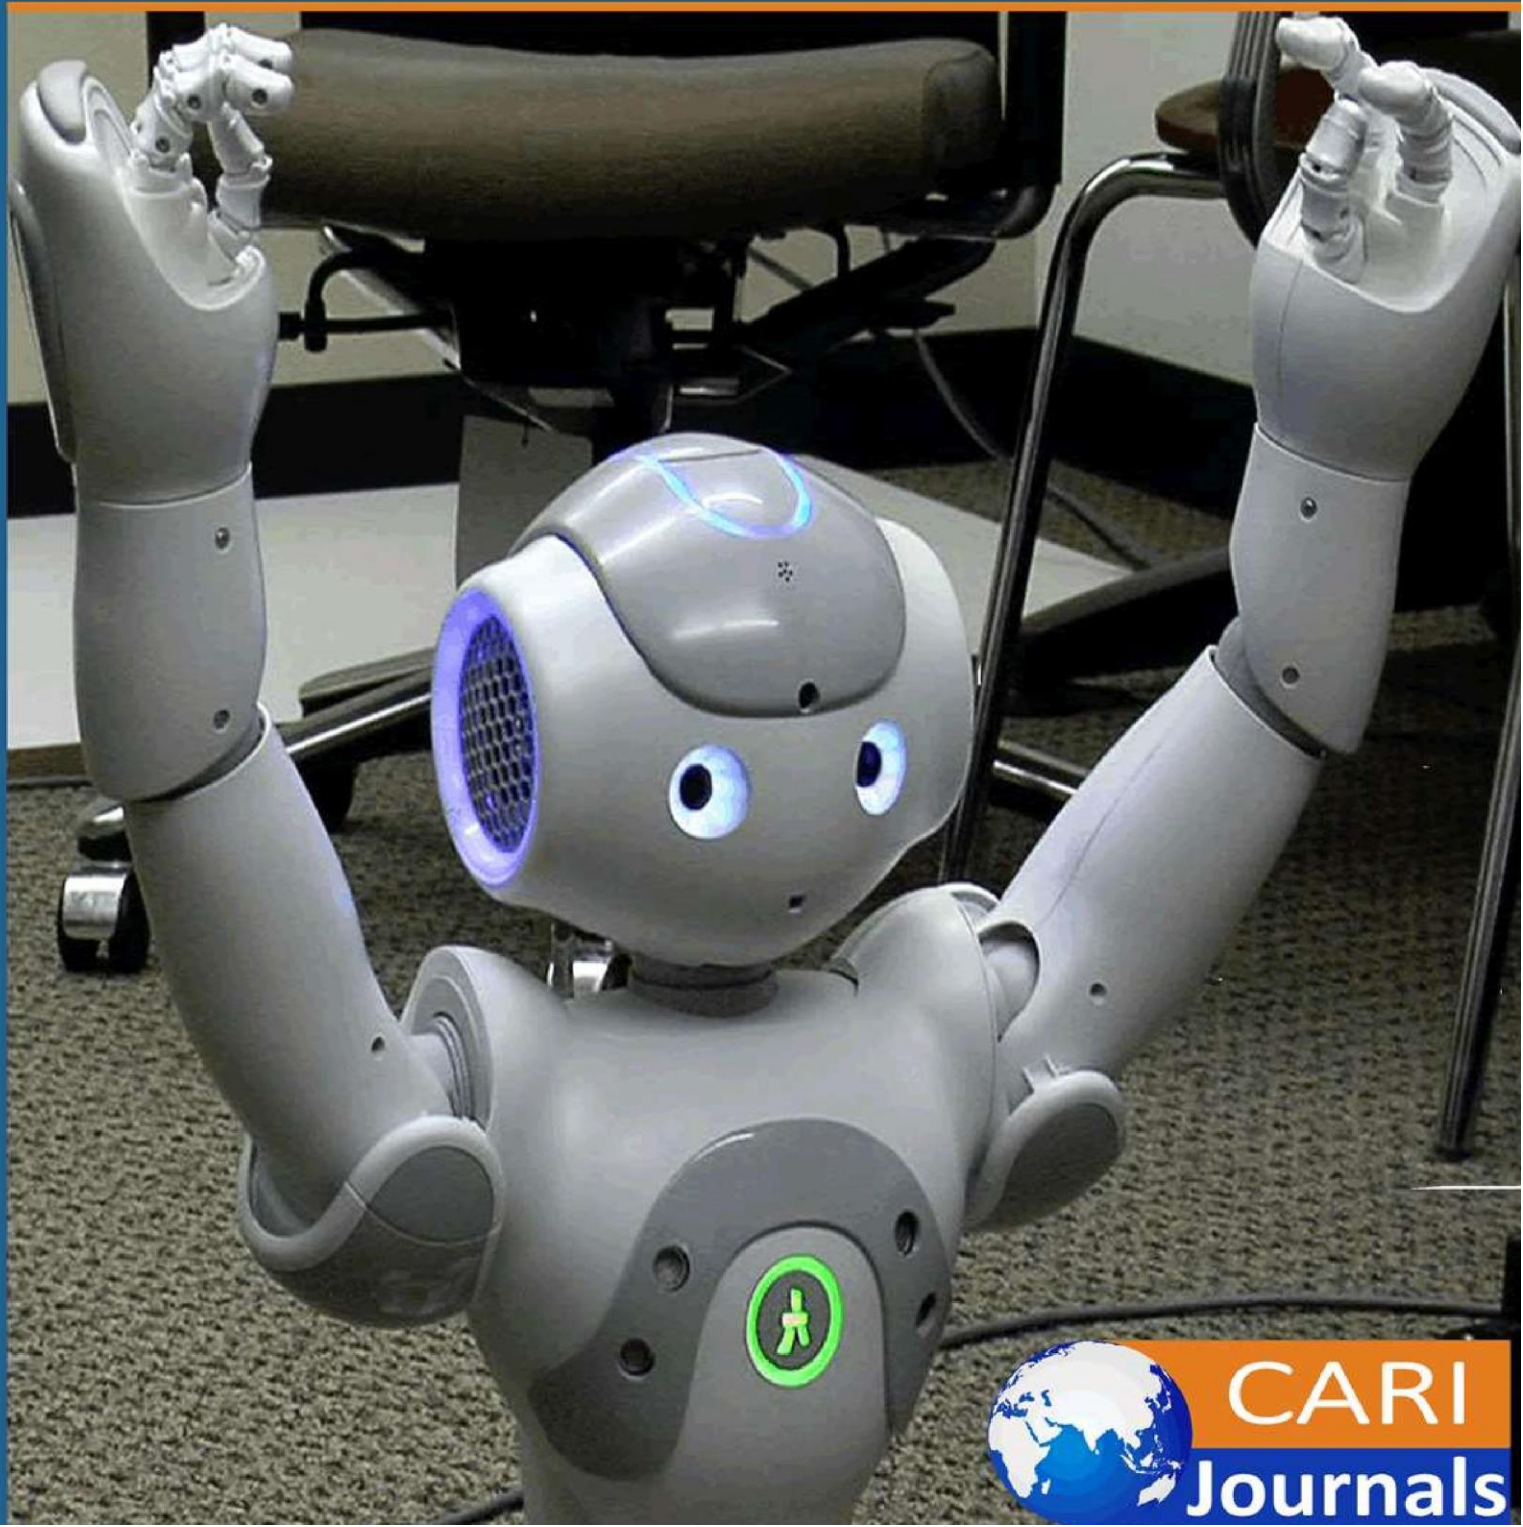

CARI Journals

International Journal of Computing and Engineering
ISSN 2958-7425 (online)
Vol. 7, Issue No. 10, pp. 48 - 59, 2025
www.carijournals.org

Federated Learning for Ransomware-Resilient Industrial IoT: A Decentralized Framework for Secure AI at the Manufacturing Edge

Aishwarya Natarajan
Amazon Web Services, USA
https://orcid.org/0009-0000-4425-4211

Accepted: 28th June, 2025, Received in Revised Form: 5th July, 2025, Published: 16th July, 2025

Abstract

Industrial Internet of Things (IIoT) deployments are facing increasing cybersecurity threats, especially with ransomware attacks on operational technology infrastructure. Traditional centralized machine learning configurations with the storage of manufacturing data in a single repository expand the attack surface area. Federated learning presents a completely new approach to conducting distributed model training across manufacturing sites with data locality. The federated learning framework uses secure aggregation protocols and encrypted communication channels to deliver intelligent systems without sending raw operational data externally. The federated model decreases the threat of ransomware propagation and exfiltration of operational data by establishing strong access control measures at the edge nodes and employing homomorphic encryption techniques. The federated approach is particularly useful in multi-site manufacturing use cases where regulatory compliance and maintaining intellectual property remain primary concerns. Demonstrations and deployments of the proposed framework in actual research problems spanning predictive maintenance, quality control, and process optimization show the model can maintain model accuracy while enhancing the operational resilience of IIoT applications. The intersection of distributed intelligence principles and cybersecurity principles provides a pathway for trustworthy AI systems in critical industrial infrastructures.

Keywords: *Federated Learning, Industrial Iot, Ransomware Resilience, Edge Computing, Privacy-Preserving AI*

International Journal of Computing and Engineering
ISSN 2958-7425 (online)
Vol. 7, Issue No. 10, pp. 48 - 59, 2025
www.carijournals.org

1. Introduction

1.1. Transformation of Traditional Manufacturing Through Connected Industrial Systems

Manufacturing paradigms have undergone a radical transformation with the emergence of cyber-physical production systems integrating heterogeneous devices, protocols, and computational resources. Contemporary industrial facilities deploy extensive sensor networks alongside programmable logic controllers, establishing multi-layered communication infrastructures between shop floor equipment and enterprise resource planning systems. Digital twin implementations enable synchronized virtual-physical asset representations, supporting condition-based monitoring and prescriptive maintenance strategies [1]. What nobody anticipated was how connecting previously air-gapped systems would turn every Ethernet port into a potential entry point. Manufacturing floors designed for efficiency now struggle with security challenges their architects never imagined.

1.2. Rising Cybersecurity Threats in Operational Technology Environments

Operational technology environments face escalating cyber risks as threat actors develop specialized capabilities targeting industrial control systems and manufacturing infrastructure. Sophisticated adversarial groups employ reconnaissance techniques, identifying vulnerable human-machine interfaces, exploiting protocol weaknesses inherent in legacy supervisory control and data acquisition architectures [2]. The latest generation of industrial malware shows disturbing sophistication. Take Triton - it learned to speak the proprietary TriStation protocol just to mess with Schneider Electric safety systems. Or consider how NotPetya turned Modbus commands into weapons, making PLCs think 200°C was actually 20°C. Field engineers tell horror stories about finding rootkits buried in HMI workstations that survived complete OS reinstalls. Once malware infiltrates an OT network, containment becomes nearly impossible - it jumps between unpatched Windows XP boxes running SCADA software, hides in firmware nobody checks, and keeps reinfecting from compromised vendor laptops.

Table 1: Evolution of Industrial Connectivity Challenges [1, 2]

Era	Primary Systems	Security Approach	Major Vulnerabilities
Pre-2000	Isolated PLCs, SCADA	Air gaps	Physical access only
2000-2010	Networked OT systems	Firewalls, VLANs	Unpatched systems, default credentials
2010-2020	IT-OT convergence	Defense in depth	Lateral movement, supply chain attacks
2020-Present	IIoT, Digital Twins [1]	Zero trust, encryption	Ransomware, APTs targeting OT [2]

1.3. Limitations of Centralized AI Approaches in Industrial Settings

Most industrial AI systems today pull data from sensors, controllers, and quality systems into central databases for analysis. This traditional approach works well in theory, but creates serious practical problems. Moving massive amounts of real-time data across networks causes delays that make split-second decisions impossible. When factories generate terabytes of vibration data or thermal images daily, network pipes simply cannot handle the load. Companies with plants in Germany, China, and Mexico have discovered that their AI ambitions have crashed into GDPR, cybersecurity laws, and data localization requirements. Worse yet, that central data lake holding twenty years of process optimization becomes a goldmine for industrial espionage. One breach exposes everything: temperature curves for heat treatment, pressure settings for injection molding, chemical ratios for proprietary coatings.

1.4. Research Objectives and Paper Organization

Federated learning offers a fundamentally different approach to industrial AI that keeps sensitive data where it belongs - at the edge. Rather than moving data to algorithms, this framework brings algorithms to data, training models locally at each factory while sharing only encrypted updates. This investigation explores how manufacturers can build intelligent systems that resist ransomware attacks and protect trade secrets without sacrificing performance. The article proceeds as follows: Section 2 examines specific attack patterns targeting industrial networks, Section 3 explains the mechanics of federated learning in factory settings, Section 4 describes cryptographic techniques for securing model updates, Section 5 showcases real deployments in predictive maintenance and quality control, and Section 6 discusses future opportunities for enhancing industrial security through distributed AI.

2. Threat Landscape and Vulnerabilities in Industrial IoT

2.1. Ransomware Attack Vectors in OT/IoT Infrastructure

Industrial control systems present unique attack surfaces that differ fundamentally from traditional IT environments. Legacy protocols like DNP3 and IEC 61850 were designed when air gaps provided security, leaving them defenseless against modern threats [3]. Ransomware groups now target the weakest links: remote access solutions hastily deployed for pandemic support, unmanaged IoT sensors with hardcoded credentials, and forgotten maintenance ports on programmable automation controllers. The Colonial Pipeline incident proved that attackers don't need to understand process control - they just need to encrypt the billing system to shut down operations. Smart factories multiply these risks exponentially, with thousands of edge devices running outdated firmware, each one a potential foothold for lateral movement into critical production systems.

2.2. Data Privacy Challenges in Manufacturing Environments

Manufacturing data reveals far more than production metrics - it exposes competitive advantages encoded in cycle times, quality parameters, and resource utilization patterns. Industry 4.0 initiatives generate unprecedented data volumes from connected machines, creating privacy risks that extend beyond traditional concerns [4]. A single vibration sensor monitoring a CNC spindle captures enough information to reverse-engineer cutting parameters and tool paths. Camera systems installed for quality inspection inadvertently record worker movements, raising surveillance concerns. Cloud-connected predictive maintenance platforms aggregate failure patterns across customers, potentially leaking one manufacturer's operational weaknesses to competitors using the same service. The distributed nature of modern supply chains compounds these risks, as tier suppliers gain visibility into OEM production schedules through shared planning systems.

2.3. Regulatory Compliance Requirements (GDPR, Industry-Specific Standards)

Manufacturing organizations navigate a maze of overlapping regulations that were never designed for connected factories. GDPR treats machine-generated data containing operator IDs as personal information, requiring consent mechanisms for systems that predate smartphones. Medical device manufacturers must satisfy FDA cybersecurity guidance while maintaining IEC 62443 compliance for industrial automation security. Automotive suppliers juggle ISO/SAE 21434 requirements for vehicle cybersecurity with TISAX audits for protecting OEM data. Cross-border data flows essential for global production coordination collide with data localization laws in China, Russia, and India. Compliance teams struggle to map IT-centric frameworks onto OT environments where a software update requires months of validation and production downtime costs millions per hour.

2.4. Case Studies of Recent Industrial Cyber Incidents

Real-world breaches reveal how theoretical vulnerabilities translate into operational disasters. The aluminum producer Norsk Hydro lost weeks of production when LockerGoga ransomware forced them to switch entire plants to manual operation. Honda's global operations ground to a halt when Snake ransomware spread through their internal network, disrupting just-in-time manufacturing across multiple continents. A German steel mill suffered physical damage when attackers manipulated furnace controls, preventing proper shutdown sequences. Closer examination shows these incidents share common patterns: initial compromise through IT systems, patient reconnaissance to understand production dependencies, and strikes timed for maximum impact during critical production runs. Recovery takes months, not days, as companies rebuild not just data but trust in their control systems.

3. Federated Learning Fundamentals for Industrial Applications

3.1. Distributed Machine Learning Architecture Overview

Federated learning flips traditional machine learning on its head by keeping data where it lives while moving algorithms to the edge. Industrial implementations leverage distributed computing frameworks that coordinate model training across heterogeneous hardware, from resource-constrained PLCs to powerful edge servers [5]. Picture an automotive manufacturer with stamping plants across three continents: instead of shipping terabytes of acoustic data to headquarters, each facility trains a local model on its press brake signatures. Think of it as a neural network that's been shattered into a thousand pieces, with each fragment learning from its data. You've got Raspberry Pis bolted to injection molding machines, and you're trying to train models alongside beefy edge servers in the QC lab. The whole thing runs like a jazz ensemble - nobody has the full score, but somehow they make music together. Half the nodes drop offline during shift changes, and others choke on memory when processing high-resolution thermal images. McMahan's team at Google cracked the problem back in 2017 - train where the data lives, share only the weight updates. Easy on paper, brutal in practice. That stamping press controller barely runs Python, while the vision inspection rig next door sports datacenter-grade hardware. Some run ancient CUDA versions, others can't even spell GPU. The aggregation algorithms don't care. FedAvg just takes whatever gradients show up and averages them, even if half the fleet is offline for maintenance. Real deployments look nothing like the papers - you get partial updates, corrupted transmissions, and that one edge device in Building C that keeps sending gradients from last Tuesday.

3.2. Local Model Training and Gradient Aggregation Mechanisms

Factory floors have morphed into distributed computing clusters where welding robots crunch numbers between spot welds. Modern implementations optimize gradient compression and sparsification to minimize bandwidth consumption, which is critical when a single robot cell generates gigabytes of trajectory data hourly [6]. The aggregation process resembles a carefully choreographed dance: edge devices compute gradients using their local datasets, apply differential privacy noise to prevent information leakage, compress updates using techniques like top-k sparsification, then transmit only these lightweight updates to the aggregation server. The aggregation server sees nothing but encrypted math - imagine trying to figure out someone's recipe by looking at their grocery bill totals. Smart factories layer this like an onion: the welding robots on Line 3 send their learnings to the plant server, which mushes together updates from all production lines. That plant server then talks to servers in Detroit, Stuttgart, and Shanghai, creating a global model without any plant knowing what the others are doing. Toyota pioneered this approach after realizing their Kaizen improvements at one plant could benefit others without revealing trade secrets. Each layer adds privacy while solving the practical problem of trying to coordinate thousands of devices over flaky industrial networks.

3.3. Communication Protocols for Edge-to-Cloud Coordination

Factory networks are a nightmare - ask any IT admin who's tried to push Windows updates through a plant firewall. You're dealing with 20-year-old switches, air-gapped segments, and network rules written when dial-up was fast. Federated learning has to play nice with this chaos. That's why everyone uses MQTT these days - it's the cockroach of protocols, surviving where fancier options die. Edge nodes queue up their model updates like teenagers texting at 3 am, waiting for that sweet spot when the network isn't clogged with production data. Protocol buffers serialize model updates efficiently, while TLS encryption protects gradients in transit. Smart factories implement priority queuing to ensure federated learning traffic never interferes with real-time control loops - a gradient update can wait, but a safety shutdown command cannot. Adaptive communication schedules adjust to production patterns, increasing update frequency during maintenance windows when networks are less congested. Some deployments leverage 5G network slicing to guarantee bandwidth for model synchronization without impacting operational systems.

3.4. Comparison with Traditional Centralized Learning Approaches

Centralized learning works beautifully in research papers but crashes hard against industrial reality. Traditional approaches require funneling all data to central servers - imagine streaming every vibration reading from every bearing in a thousand-machine factory to the cloud. Federated learning keeps data local, trains models where the sensors live, and shares only mathematical updates. Centralized systems offer perfect synchronization and simple debugging but create single points of failure and massive attack surfaces. Federated approaches trade some model convergence speed for radical improvements in privacy, reduced bandwidth costs, and regulatory compliance. While centralized training might achieve slightly better accuracy in controlled conditions, federated learning wins in the real world where network outages happen, data sovereignty matters, and a compromise can't expose decades of production secrets. The distributed approach also enables continuous learning - models improve even when plants lose internet connectivity, synchronizing improvements when connections are restored.

Table 2: Comparison of Centralized vs Federated Learning in Industrial Settings [5, 6]

Aspect	Centralized Learning	Federated Learning
Data Location	Cloud/Central servers	Edge devices
Network Requirements	High bandwidth, constant connectivity	Intermittent, low bandwidth [6]
Privacy Protection	Data leaves the premises	Data stays local
Latency	High (cloud round-trip)	Low (edge inference)
Failure Impact	System-wide outage	Isolated node failures
Regulatory Compliance	Complex for multi-jurisdiction	Simplified - data doesn't cross borders
Attack Surface	Single high-value target	Distributed, harder to compromise fully
Scalability	Vertical (bigger servers)	Horizontal (more edge nodes) [5]

4. Security-Enhanced Federated Learning Framework

4.1. Secure Aggregation Protocols for Model Updates

Getting factories to share model improvements without revealing trade secrets requires cryptographic gymnastics that would make a blockchain developer sweat. The latest protocols use homomorphic encryption to let servers add encrypted gradients without decrypting them - like doing math while wearing a blindfold [7]. Picture this: a bearing manufacturer's edge device encrypts its failure prediction gradients, sends them to an aggregation server that combines updates from competitors' plants, and produces an improved global model without anyone learning that Plant A discovered vibration patterns predicting failure weeks before anyone else. Real implementations layer multiple techniques: secret sharing splits each gradient across multiple servers, differential privacy adds just enough noise to hide individual contributions, and zero-knowledge proofs verify computations without revealing inputs. BMW's production network reportedly uses five-party computation, where no three servers can reconstruct the original data together.

4.2. End-to-End Encrypted Communication Channels

Factory networks weren't built for cryptography - they were built for speed and reliability when Reagan was president. Retrofitting end-to-end encryption onto industrial systems requires careful engineering to avoid breaking real-time guarantees [8]. Modern deployments establish TLS tunnels from edge devices to aggregation servers, but that's just the start. Keys need rotation without disrupting production, certificate authorities must work across air-gapped networks, and hardware security modules protect root keys from physical tampering. The tricky part comes when dealing with resource-constrained devices - that temperature sensor running on an 8-bit microcontroller can't handle AES-256. So you end up with this patchwork of crypto: ARM devices

get ChaCha20 because it's fast without hardware acceleration, while the big boys run AES-GCM. Key rotation is a mess - you need perfect forward secrecy so next month's breach doesn't expose last year's production data. Try explaining to plant managers why their 15-year-old HMI needs a firmware update just for new crypto libraries.

Table 3: Security Mechanisms for Industrial Federated Learning [7, 8]

Security Layer	Technology	Purpose	Industrial Constraints
Gradient Protection	Homomorphic encryption [7]	Hide model updates	Computational overhead on edge devices
Communication	TLS 1.3, ChaCha20 [8]	Encrypt transmissions	Legacy device compatibility
Authentication	mTLS, Hardware TPM	Verify edge nodes	Certificate management at scale
Aggregation	Secure multi-party computation [7]	Private model merging	Requires multiple trusted servers
Anomaly Detection	Statistical monitoring	Detect poisoning attacks	Must distinguish from legitimate variance

4.3. Robust Access Control and Authentication at Edge Nodes

Securing thousands of edge devices scattered across a factory floor makes enterprise IT look like child's play. You can't just slap Active Directory on a PLC and call it done. Industrial federated learning systems implement defense in depth: hardware-based device identity using TPM chips, mutual TLS authentication for every connection, and role-based access control that understands the difference between a maintenance technician and a data scientist. Certificate management becomes nightmarish at scale - imagine provisioning unique identities for every sensor in a facility, then doing it again across fifty plants. Smart implementations use intermediate certificate authorities at each site, letting local teams manage device identities while maintaining global trust chains. Time-based access tokens expire after shifts end, and geo-fencing ensures that edge devices can only connect from expected locations.

4.4. Anomaly Detection and Adversarial Attack Mitigation

Poisoning attacks against federated learning aren't theoretical - they're Tuesday morning for security teams. Malicious nodes can submit crafted gradients designed to corrupt the global model, making it misclassify defects or ignore safety warnings. Detection systems monitor statistical properties of incoming updates: sudden spikes in gradient magnitudes, updates that consistently point opposite to the consensus, or nodes whose contributions degrade model performance. Byzantine-robust aggregation algorithms like Krum or trimmed mean throw out suspicious updates before they poison the well. But clever attackers adapt - they'll submit normal updates for weeks to build a reputation, then strike during critical production runs. Advanced defenses use ensemble methods, maintaining multiple global models trained on different subsets of nodes and comparing

their predictions. When the welding robot in Bay 7 starts sending updates that make every model except one predict nonsense, you know something's wrong.

5. Implementation Strategies and Industrial Use Cases

5.1. Architecture Design for Multi-Site Manufacturing Networks

Designing federated learning for a single factory is like teaching your dog to sit. Scaling it across global manufacturing networks is like conducting a symphony where half the orchestra is on Mars. Knowledge-based architectures for smart manufacturing networks provide blueprints, but reality hits differently when your Shanghai plant runs three shifts while Detroit is closed for Thanksgiving [9]. The architecture that works starts with regional hubs - Asian plants federate together during their day shift, then pass the baton to European facilities, creating a follow-the-sun model of training. Volkswagen allegedly structures its network with plant-level aggregators feeding into regional clusters (Americas, Europe, Asia), with a global orchestrator that only touches metadata. The clever bit is using production schedules to trigger training cycles - when Line 3 switches from sedans to SUVs, it triggers local retraining that propagates improvements globally without revealing what model is being built where.

5.2. Scalability Considerations for Global Factory Deployments

Manufacturing scalability looks clean in architecture diagrams, but turns into a street fight when deployed globally [10]. You start with ten edge devices in a pilot program, and everything works great. Scale to ten thousand across fifty plants, and suddenly you're drowning in certificate renewals, firmware updates, and that one facility in Brazil that keeps disconnecting every Tuesday. The math says federated learning scales linearly with nodes. The reality is that your aggregation server melts when a thousand devices try to upload gradients simultaneously after a power outage. Smart deployments use hierarchical aggregation with exponential backoff - think of it as crowd control for robots. Geographic distribution helps: Asian plants aggregate locally before talking to global servers, reducing transcontinental traffic. But the real killer is heterogeneity - mixing ancient Siemens controllers with cutting-edge NVIDIA boards means your slowest node determines your training speed.

Table 4: Industrial Federated Learning Deployment Patterns [9, 10]

Scale	Architecture	Aggregation Strategy	Key Challenges
Single Plant	Flat hierarchy	Direct to the plant server	Device heterogeneity
Multi-Plant (Regional)	Two-tier	Plant → Regional → Global	Time zone coordination
Global Enterprise	Three-tier hierarchical [9]	Follow-the-sun aggregation	Cross-border regulations
Supply Chain Network	Mesh with trusted clusters	Selective sharing by tier [10]	Trust establishment between competitors

5.3. Performance Evaluation: Accuracy, Latency, and Resource Utilization

Measuring federated learning performance in factories is like judging a beauty contest in the dark. Academic papers obsess over model accuracy, but plant managers care about different metrics: Will it catch defects before they ship? Can it run without slowing production? What happens when half the edge nodes lose power? Real deployments track convergence speed across heterogeneous hardware, measuring how long before a model trained in Germany helps prevent defects in Mexico. Latency matters differently here - a centralized model might be slightly more accurate. Still, if it takes three seconds to classify a defect on a line moving at sixty parts per minute, you've already shipped a bad product. Resource utilization gets tricky with legacy hardware: Allen-Bradley PLC has spare cycles between control loops, but when you touch its memory allocation, you'll trigger safety shutdowns. Ford reportedly maintains shadow models to compare federated versus centralized performance, finding federated learning catches location-specific defects that global models miss entirely.

5.4. Real-World Applications: Predictive Maintenance, Quality Control, and Process Optimization

Federated learning shines brightest where local conditions matter. Take bearing failure prediction - a bearing in Phoenix fails differently than one in Michigan winters, but both patterns help predict failures in Mexico City. Bosch's predictive maintenance system reportedly federates vibration models across automotive plants, with each facility contributing patterns from their specific equipment mix without revealing which car models they produce. Quality control gets interesting when camera systems at different plants learn collaboratively - scratches look different under German LED lighting versus Chinese fluorescents, but federated models adapt to both. Process optimization pushes boundaries further: injection molding parameters that work in humid Singapore fail in dry Arizona, but federated learning finds the underlying relationships. The killer app might be energy optimization - factories share patterns of equipment scheduling and HVAC

usage without revealing production volumes, collectively reducing energy consumption while keeping competitive intelligence secret.

Conclusion

Federated learning turns industrial cybersecurity from a defense to an opportunity. By keeping sensitive manufacturing data on the edge and enabling collective intelligence, factories can enjoy both AI capabilities while protecting their crown jewels from ransomware gangs or competitors. This technology is now beyond the 'interesting project' stage of academia into systems that can be deployed in the world's untidy realities of global manufacturing, such as old PLCs feeding into modern edge servers, network reliability to all model gym failures, as well as the regulatory complexities of country borders. With 'success' stories in diverse sectors including automotive and aerospace, federated learning provides benefits such as defect detection within distributed learning that centralised models do not, accurate failure predictions that use the noise of their local conditions, and optimisation of processes while preserving the sanctity of existing intellectual property. As manufacturing networks become ever more connected and as increasingly sophisticated threats develop, federated learning offers a way to improve both intelligence and resilience. The next challenge is to extend these systems to the domain of real-time control, increasing the decision cycle time to milliseconds on the edge and providing cryptographic assurances. Smart factories of the future, if they are ever to fulfil their destiny, will train their models where their data reside, share insights without disclosing secrets, and build collective intelligence that makes the whole more decidedly greater than the sum of its parts.

References

[1] Rajesh Kumar Dhanaraj, "Digital Twins in Industrial Production and Smart Manufacturing: An Understanding of Principles, Enhancers, and Obstacles," IEEE Xplore (Wiley-IEEE Press), 2024. [Online]. Available: https://ieeexplore.ieee.org/book/10705130. [Accessed: 2024].

[2] Mukund Bhole, et al., "Enhancing Industrial Cybersecurity: Insights From Analyzing Threat Groups and Strategies in Operational Technology Environments," IEEE Open Journal of the Industrial Electronics Society, January 2025. [Online]. Available: https://repositum.tuwien.at/bitstream/20.500.12708/212806/3/Bhole-2025-IEEE%20Open%20Journal%20of%20the%20Industrial%20Electronics%20Society-vor.pdf. [Accessed: January 2025].

[3] Soobia Saeed, et al., "Ransomware: A Framework for Security Challenges in Internet of Things," IEEE Conference Publication (ICCIS), 24 November 2020. [Online]. Available: https://ieeexplore.ieee.org/document/9257660/references#references. [Accessed: 24 November 2020].

[4] Md Mehedi Hassan ONIK, et al., "Personal Data Privacy Challenges of the Fourth Industrial Revolution," IEEE Conference Publication (ICACT), 02 May 2019. [Online]. Available: https://ieeexplore.ieee.org/document/8701932. [Accessed: 02 May 2019].

[5] Fan Meng, et al., "An Overview of PAI: Distributed Machine Learning Platform," IEEE Conference Publication (IMCEC), 26 January 2023. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10020005/figures#figures. [Accessed: 26 January 2023].

[6] Hasin Us Sami, Basak Guler, "Secure Gradient Aggregation with Sparsification for Resource-Limited Federated Learning," IEEE Transactions on Communications, 2024. [Online]. Available: https://basakguler.github.io/SG_TCom_2024_sparse.pdf. [Accessed: 2024].

[7] Tamer Eltaras, et al., "Efficient Verifiable Protocol for Privacy-Preserving Aggregation in Federated Learning," IEEE Xplore, 2023. [Online]. Available: https://ieeexplore.ieee.org/stampPDF/getPDF.jsp?arnumber=10121168. [Accessed: 2023].

[8] Rolf Oppliger, "End-to-End Encrypted Messaging," Artech eBooks - IEEE Xplore, 2020. [Online]. Available: https://ieeexplore.ieee.org/book/9118792. [Accessed: 2020].

[9] Michael P. Papazoglou, et al., "A Reference Architecture and Knowledge-Based Structures for Smart Manufacturing Networks," IEEE Software, 23 April 2015. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/7093038.

[10] G. Putnik, et al., "Scalability in Manufacturing Systems Design and Operation: State-of-the-Art and Future Developments Roadmap," CIRP Annals - Manufacturing Technology, 2013. [Online]. Available: https://ykoren.engin.umich.edu/wp-content/uploads/sites/122/2014/05/Scalability-2013.pdf.