TOPIC GROUP MODULAR PRODUCTION

Reference architecture

Manufacturing-X

Guidance Board

Contents

IMPRINT

Publisher

Manufacturing-X Guidance Board

Project Office Manufacturing-X

Open Industry 4.0 Alliance Implementation GmbH

Willy-Brandt-Platz 6

81829 Munich, Germany

Version History

[June 23, 2026 - Version 1.0 (First Release)]

Authors

Christian Methe (twindustrial GmbH)

Haris Avgoustinos (KIWAI.EXPO21XX)

Pascal Rübel (SmartFactory-KL)

Thomas Tauchnitz (Tautomation GmbH)

Co-Autoren / Topic-Group-Teilnehmende

Andreas Bamberg (Merck KGaA)

Sebastian Bartsch (YARDSTICK Robotics GmbH)

Richard Bormann (Fraunhofer-Gesellschaft e.V. / Fraunhofer IPA)

Domenico D'Agostino (YARDSTICK Robotics GmbH)

Florian Düwel (Fraunhofer-Gesellschaft e.V. / Fraunhofer IOSB)

Andreas Ebner (Fraunhofer-Gesellschaft e.V. / Fraunhofer IOSB)

Rüdiger Fritz (Aktorion Gemini Solutions ENK)

Sebastian Haertner (Merck KGaA)

Christoph Hock (Open Industry 4.0 Alliance)

Anselm Klose (FLEXCELERATE Solutions GmbH)

Prof. Dr. Frank Köster (Deutsches Zentrum für Luft- und Raumfahrt e.V.)

Markus Lieret (Siemens AG)

Anna Menschner (SEMODIA GmbH)

Anja Muesch (Merck KGaA)

Felix Spenrath (Fraunhofer-Gesellschaft e.V. / Fraunhofer IPA)

Robin Taba (Deutsches Zentrum für Luft- und Raumfahrt e.V.)

1 Introduction

This document is intended for companies, research and development consortia, standardization initiatives, and decision-makers involved in shaping the future of flexible, connected, and interoperable production systems. It specifically addresses stakeholders seeking to evaluate, develop, or implement modular production approaches in industrial applications. These include industrial users, machine and plant manufacturers, software providers, integrators, research institutions, and funded projects within and beyond Manufacturing-X.

The “Modular Production” use case was launched as part of the Catena-X research project as early as July 2021. In February 2024, the use case was also taken up in Factory-X. Subsequently, further ecosystems addressing Modular Production emerged within the Manufacturing-X funding framework. In the same year, the Manufacturing-X Guidance Board, as the central coordinating body within the Manufacturing-X initiative, decided to establish the Modular Production Topic Group.

Manufacturing-X is an initiative of Plattform Industrie 4.0, industry, policymakers, and the scientific community and provides the overarching framework. The objective of Manufacturing-X is to establish industrial data ecosystems in which companies can use data in a sovereign, interoperable, and value-oriented manner. This goes beyond the exchange of individual data points and encompasses the digital integration of value chains, production systems, machines, products, and services. In this context, Modular Production is a key application area because it requires standardized descriptions, interoperable interfaces, digital twins, semantic models, and trusted data spaces.

The central question is how Modular Production can be understood across industries, positioned from a technical perspective, and made economically viable. This document describes the initial situation, the motivation of the participating Manufacturing-X projects, key terms and common questions, as well as the reference architecture developed by the group. It also presents exemplary use cases from different industrial sectors. These illustrate the requirements, solution approaches, and potential benefits associated with Modular Production and highlight common patterns that can be identified across industry boundaries.

The relevance of Modular Production is increasing significantly. Industrial production is under growing pressure to respond more quickly to volatile markets, shorter product life cycles, increasing product variety, unstable supply chains, skilled labor shortages, and higher sustainability requirements. Under these conditions, traditionally rigid production structures are increasingly reaching their economic and technical limits. Modular production concepts provide a structural response: production resources, processes, and digital services are described, encapsulated, and combined in ways that allow them to be deployed more flexibly, reconfigured more quickly, and integrated more effectively across company or site boundaries.

Modular Production was already addressed as a dedicated use case in Catena-X, the first industrial data ecosystem. Under the use-case leadership of Fraunhofer IPA, key foundations were developed and later documented and made available for further use in Tractus-X as the "Modular Production KIT." The KIT aims to increase the flexibility and reliability of industrial production through standardized services, interfaces, and data models. Key results include the Shop Floor Information Service (SIS), with its Production Forecasting and Production Tracking components; the use of common semantic models for production and process information; and the Bill of Process (BoP) as a standardized description of manufacturing processes, process steps, dependencies, and process parameters. These developments laid important groundwork for the cross-company exchange of production planning data, process parameters, and traceability information within the data space. The official documentation of the "Modular Production KIT" is available on the Eclipse Tractus-X website: https://eclipse-tractusx.github.io/docs-kits/category/modular-production-kit/

Factory-X is one of the central Manufacturing-X projects with a particular focus on industrial production. Its emphasis is on end-to-end, scalable, and vendor-independent solutions for the connected factory. Flexible production processes, automated resource selection, product changeovers, digital capabilities of equipment, and integration via data spaces are key building blocks. The concepts developed in Factory-X make important contributions to the discussion around Modular Production, particularly where production resources are no longer rigidly assigned to a specific process but are instead used dynamically based on their capabilities.

The Manufacturing-X Guidance Board plays a coordinating role within the initiative. It provides orientation, consolidates cross-cutting topics, and supports exchange between the individual Manufacturing-X projects. Topic Groups were established within this framework. Their purpose is to jointly address cross-cutting topics relevant to multiple projects, exchange experiences, identify overlaps, and develop potential common foundations.

The Modular Production Topic Group was established because several Manufacturing-X projects independently address similar questions. Across different industries, concepts are emerging for more flexible production, modular descriptions of equipment, the use of digital twins, automated configuration of production resources, and data-space-enabled collaboration. Without coordination, there is a risk of parallel and incompatible developments. At the same time, the diversity of the projects offers an opportunity to identify common principles, reusable architectural building blocks, and terminology that can be applied across industries.

The Topic Group includes participants from Factory-X, Robot-X, RO-X, Process-X, TWIN4TRUCKS, and Semiconductor-X, among others. The projects contribute different industrial perspectives, ranging from discrete manufacturing and robotics to commercial vehicle production, the process industry, and semiconductor manufacturing. Factory-X approaches Modular Production primarily from the perspective of flexible, self-optimizing production processes and vendor-independent integration. Robot-X and RO-X address modular and AI-supported robotics solutions as well as their development, commissioning, and operation. Process-X contributes the perspective of modular process plants with standardized descriptions such as MTP, PEAs, and POL. TWIN4TRUCKS focuses on digital twins, production logistics, quality assurance, and adaptive planning in commercial vehicle production. Semiconductor-X broadens the perspective to include cross-factory production and supply-chain relationships, semantic interoperability, and data-driven optimization of complex interactions between materials, processes, tools, and products.

At the beginning of the collaboration, several fundamental questions were at the forefront. First, it was necessary to clarify what different industries understand by Modular Production and which terminology, degrees of freedom, and target visions they use. The different levels of maturity were equally relevant: while established approaches to modular automation such as MTP already exist in the process industry, other sectors are still at a more conceptual or prototypical stage. The group also examined the economic improvements expected in the respective industries, for example through reduced setup times, faster reconfiguration, better resource utilization, greater resilience, lower integration efforts, or shorter time-to-market.

Another focus was the question of whether common patterns can be identified despite different industry-specific logics. The discussions showed that many of the challenges are similar: production resources need to be described unambiguously, capabilities need to be available in machine-readable form, digital twins require semantic structure, data needs to be usable across company boundaries, and software artifacts should be reusable wherever possible. This also raises the question of whether the individual projects can provide the basis for a cross-industry definition of Modular Production, a common reference architecture, and a long-term expert group capable of continuing this work.

This document is not intended to serve as a final standard. It summarizes the current state of the Topic Group's joint work, provides context for key terms and solution approaches, and describes a reference architecture as guidance for further discussion and implementation. The objective is to improve interoperability and alignment between the projects, avoid duplicate developments, and establish a common foundation for the further development of Modular Production within Manufacturing-X and beyond.

2 Research Projects and Focus Areas

At the start of the Topic Group in 2024, almost all funded and non-funded X-projects were asked about their interest in participating. Some research projects focused on other areas, such as supply chain optimization or Digital Product Passport approaches, meaning that production itself was not part of all respective development projects. However, this does not mean that Modular Production is not relevant to these industries.

X-Projekt	Focus of Modular Production
FACTORY-X	Flexible, self-optimizing production processes through product changeovers and automated resource selection. The goal is resilient, responsive production with decentralized data storage and vendor-independent integration.

ROBOT-X	AI-supported robotics solutions for flexible, modular production. Development of a decentralized data and service ecosystem for faster development, commissioning, and maintenance.
ROX Enabling AI Robotics	End-to-end modularization of the robot lifecycle. Modular building-block systems, sovereign data sharing, and circular use to reduce integration times.
PROCESS-X	Standardized modules (MTP, PEAs, POLs) for flexible, reconfigurable process engineering. Plug & Produce and energy-efficient, adaptive plant configuration.
TWIN4TRUCKS	Digital twins and AI for connected commercial vehicle production. Integration of production, logistics, and quality assurance through the digital thread, with adaptive planning and worker assistance.
SX SEMICONDUCTOR-X	Cross-factory modular production and further processing of semiconductors along the entire supply chain. Use of a digital graph model, including the AAS of the material, to integrate knowledge about materials, processes, and tools. The goals are traceability, zero waste, and high utilization through data-driven control of the process chain.

3 Reference Model as an Enabler for Industrial AI

The developed reference architecture provides the foundation for seamless integration of IT and OT systems and end-to-end data connectivity from the shop floor to the connected world of data spaces. It shows how components can be intelligently encapsulated and connected to support new business models, accelerate innovation, and make plants more flexible. At the same time, the structure of the MX-Port can be mapped to the elements of the reference architecture. Existing investments should be protected and existing applications integrated. Figure 1 shows the architectural building blocks from bottom to top.

At the field level, numerous bus systems and protocols have evolved as manufacturers addressed industry-specific requirements with their own, often proprietary solutions. This makes integrating new devices into existing production environments more difficult. Traditional bus systems are based on PLCs with associated sensors and actuators. Smart field devices such as RFID sensors or robots, on the other hand, operate independently and communicate via defined interfaces of the Edge Layer. In the reference architecture, they are integrated on edge devices rather than in the PLC, while real-time and safety requirements must continue to be met. The Adapter Layer of the MX-Port concept is implemented through interface applications that transfer data to a digital OT data bus. This marks the point where traditional automation ends and data can be processed flexibly (see L1 in Figure 1).

A Smart Machine Logic Controller combines modular automation functions, known as skills, which enable flexible commissioning and reconfiguration. Instead of a device-centric approach, the focus is on functions: proprietary interfaces of PLCs, smart sensors, and actuators are encapsulated in reusable skills. Existing skills can be used to derive complex, adaptable functions without changing the control code. Hardware can also be replaced without requiring changes to the skills.

The Smart Machine Interface transforms operational OT data into a standardized format and serves as the central access point to the machine network. OPC UA for Machinery is recommended for this purpose, enabling interoperable collection of machine states, process monitoring, and semantic interpretation of data. The skills of individual components are made available through building blocks. The standardized interface enables consistent user interfaces and analytics tools, reduces familiarization time and errors, and thereby increases efficiency in the production process.

Above the IT/OT level, the entire lifecycle of assets must be considered. The Asset Administration Shell (AAS) consolidates all data points through standardized interfaces and provides a unified view of static and dynamic information. Manufacturer-specific data is converted into AAS-compliant formats and thereby integrated. Data from different sources can be aggregated in a standardized manner and used for reports that provide transparency throughout the asset lifecycle.

Both the use of OPC UA for Machinery and the use of AAS submodels represent Layer 2 of the MX-Port concept (see L2 in Figure 1).

Both variants provide their own standardized, addressable interfaces. This allows a skill representation to be accessed either directly (Orion variant) or through additional encapsulation in the AAS (Hercules variant). In both variants, the interfaces are described and made accessible in the AAS using the Asset Interfaces Description submodel. For this reason, only the AAS-based IT Data Layer is designated as Layer 3 of the MX-Port in the reference architecture (see L3 in Figure 1).

There are two different variants for connectivity above the AAS. The Leo variant of the MX-Port exchanges data directly via the AAS. Hercules and Orion use the Eclipse Data Space Connector at this level. The connector enables data to be securely integrated into existing data spaces, embedded within access rights management (see L4 in Figure 1). The discovery of other partners and their data is also implemented at this level (see L5 in Figure 1).

Semantic descriptions enable the consistent and scalable use of internal and external data for data-centric enterprise management. AI benefits from contextually described data, allowing it to identify patterns more quickly and make predictions. A consistent data structure also makes AI more precise, transparent, and easier to scale.

The IT data bus enables decentralized, autonomous units to be integrated into factory control. Intelligent agents are used for this purpose, independently pursuing objectives and cooperating with other agents. This concept increases the flexibility, scalability, and adaptability of manufacturing systems, for example in product customization or when responding to unexpected conditions.

Agents manage individual assets such as products, manufacturing cells, or transport units and are divided into service, product, and resource agents. Service agents coordinate data exchange and services via connectors and manage access according to defined policies. Product agents handle production planning. Resource agents execute orders and optimize scheduling and machine utilization.

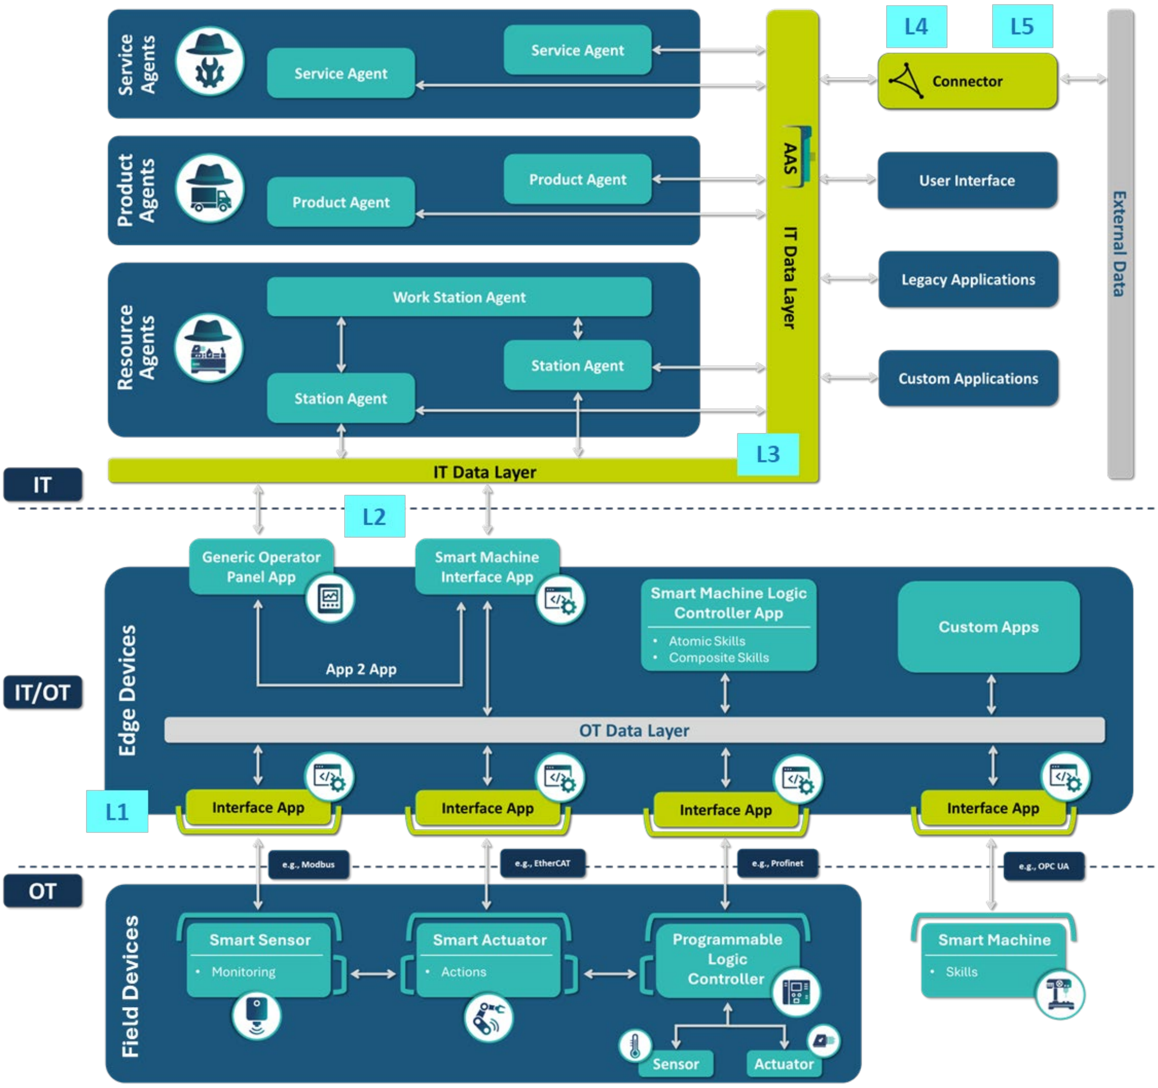

Figure 1: Reference Architecture, Version 2026/06 – 1

4 Economic Benefits

The strategic and operational success of Modular Production can be measured and managed through a multidimensional system of Key Performance Indicators (KPIs). These KPIs highlight the performance advantages compared with monolithic production concepts. In terms of production flexibility, responsiveness to market changes increases, enabling rapid adaptation to changing product portfolios and small batch sizes. This is reflected primarily in reduced Changeover Time when switching products, enabled by standardized mechanical and software-based Plug & Produce concepts. In addition, the Time to Reconfigure or setup time for entire plants or production sections can be reduced from several days to just a few hours by recombining previously validated modules. At the same time, the Number of Producible Variants increases, enabling greater product variety up to customer-specific series production on the same physical infrastructure.

The economic efficiency dimension supports optimal resource utilization and leads to lower unit costs despite high product variability. One key lever is Overall Equipment Effectiveness (OEE). Sustainable improvements in equipment availability and utilization are achieved by AI agents autonomously optimizing resource selection and workpiece routing during ongoing operations. Product Introduction Cost is also reduced, as engineering, planning, and physical conversion costs associated with introducing new product lines are lowered through the reuse of existing module structures. This is complemented by increased Energy Efficiency through targeted energy load management, in which individual modules can be activated or switched to an energy-saving mode depending on current capacity requirements or production status.

With regard to Time-to-Market, innovation cycles are shortened, accelerating the transition from product development to actual series production. Optimized Time-to-Production enables faster physical and digital commissioning, as modules can be put into operation directly via standardized, vendor-independent interfaces such as MTP. In parallel, Development Cycle Time for new production processes is reduced through the use of previously validated, certified, and operational modules.

Adaptability and resilience address high operational reliability, fault tolerance, and stable quantitative output despite external or internal disruptions. In the event of a failure, Mean Time to Recovery (MTTR) is reduced because defective modules can be replaced via Plug & Produce. A consistently high Order Fulfillment Rate ensures that customer orders can still be reliably fulfilled when individual parts of a plant fail, with agent-based routing flexibly directing products through alternative modules offering equivalent functionality. Process Adaptation Speed enables the factory to respond automatically and in real time to dynamic process changes, changing quality data, or unexpected bottlenecks in material flow. Combined with optimized Material Utilization and a reduced Waste Rate, scrap can be reduced through more precise, data-driven, cross-module process control and continuous quality analysis by AI agents.

5 Application Examples

The respective project results are described within the individual projects. The following section presents three use cases that illustrate the practical application and benefits of Modular Production in their respective industrial domains.

5.1 Use Case: Truck Production

At the Daimler Truck plant, assembly is characterized by an exceptionally high level of product variety. No two trucks are alike. Different vehicle configurations, for example with electric, diesel, or fuel-cell powertrains, require different work steps, tools, and quality inspections. Traditional assembly-line structures are increasingly reaching their limits in dealing with this level of variety, as they can only adapt to short-term product changes, different configurations, or new requirements to a limited extent.

TWIN4TRUCKS therefore aims to break up these rigid structures and develop them into modular, reconfigurable production modules. Instead of carrying out assembly operations exclusively along a fixed production line, work steps are organized in autonomous production modules orchestrated by digital twins. Such modules can, for example, be used for axle assembly, cab installation, or battery assembly for electric trucks. Depending on the order situation and vehicle configuration, these cells can be assigned dynamically.

Production is controlled through an agent system and the Digital Twin of the truck being manufactured. When determining the production sequence, the system considers which modules are capable of performing each work step, where bottlenecks or maintenance periods exist, and how energy consumption can be minimized. These decisions are based on live data, for example from UWB localization, sensors, and the production data space. The modular approach increases production flexibility because modules can be dynamically recombined in response to product variants or disruptions. From a KPI perspective, reconfiguration time in particular is optimized through greater efficiency and rapid adaptability. At the same time, routing flexibility improves because reusable modules can be better adapted to different production situations.

One specific use case within TWIN4TRUCKS addresses the traceability of load carriers between the plant and its suppliers. The underlying challenge is that data use and data exchange across company boundaries represent realistic and increasingly important requirements in industrial value chains. Today, such requirements are often implemented through individual interfaces or industry-specific standards, resulting in corresponding implementation effort. The load carrier use case examines how a physical asset with different states can be jointly managed. Both sides should have the ability to read the status of a load carrier and modify it when required.

The starting point is that special load carriers are circulating assets used in a wide variety of configurations. At the same time, the load carrier itself can be the cause of defects. Without systematic tracking, however, potentially defective load carriers cannot be clearly identified. In addition, the link between the load carrier, the transported parts, and the vehicle is often missing. The solution therefore consists of technically tracking special load carriers, linking parts, load carriers, and vehicles, and establishing communication via a Gaia-X infrastructure.

Another use case examines the provision and use of production resources and digital services via the data space, using an electric tightening tool service as an example. The objective is to make data from the electric tightening tools available and analyze it promptly in order to identify quality issues at an early stage and enable predictive maintenance. In assembly, electric tightening tools are used whose process data is currently only available locally and has to be checked manually. As a result, quality deviations and signs of wear in the tools are often detected only at a late stage, potentially leading to rework or machine downtime.

By making data from the electric tightening tools available and analyzing it promptly, torque curves can be visualized and threshold violations can be reported automatically. This supports early fault detection and enables predictive maintenance. The result is a data-based view of the tightening tools' torque curves that can be used for quality assurance and maintenance.

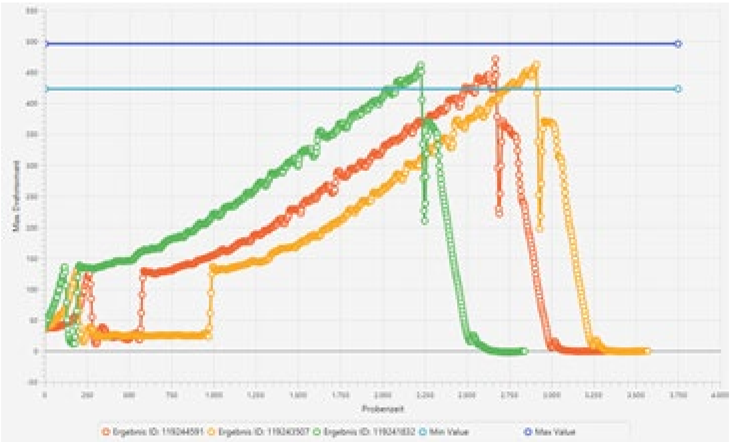

Figure 2: Visualization of Electric Tightening Tool Torque Data

5.2 Use Case: Modularization in the Process Industry

Chemical plants are optimized and built for a single product and then operated for decades – or at least, that used to be the case. Flexible markets, high product variety, and cost pressure require flexible production that can be adapted quickly and cost-effectively to capacity fluctuations or product changes. Modular plants are the solution of choice.

NAMUR, the User Association of Automation Technology in Process Industries, recognized early on that modular production also requires a modular automation architecture. As early as 2013, NAMUR Recommendation NE 148, "Automation Requirements Relating to Modularisation of Process Plants," was published. In the autumn of the same year, Wago, sponsor of the NAMUR General Meeting, proposed the definition of a "Module Type Package" (MTP), which represents the automation of modules in a standardized information model. The MTP can be compared to a device driver, as it enables modules – known as Process Equipment Assemblies (PEAs) – to be integrated into the "orchestration" of the overall plant within minutes. This concept is known as "Plug and Produce."

The MTP file contains the following information:

- The process engineering services offered by the module, e.g. “temperature control with a fixed setpoint” or “mixing two substances at a specified ratio”
- The module’s human-machine interface (HMI), including static elements such as piping and dynamic values such as current process values
- The data generated by the module, e.g. process values, and the data it requires, e.g. setpoint specifications
- Alarms and event messages generated by the module

When the module is manufactured and automated – usually using a Programmable Logic Controller (PLC) – the PLC engineering tool generates the MTP file. The PLC is then connected to the higher-level automation system. This is referred to as the Process Orchestration Layer (POL), which controls and coordinates the modules.

The knowledge about the module that the POL requires for orchestration is contained in the MTP file. Once the MTP files have been imported, the module can be monitored and operated from the POL. Process values, alarms, and event messages are displayed and evaluated by the POL system in the same way as if they originated directly from the control system of the respective POL manufacturer.

The MTP concept is already used by several companies in the process industry, particularly where speed is critical: process development in laboratories and pilot plants, rapid production of products for market launch, and fast capacity expansion in response to increasing demand. Around 20 manufacturers of process engineering machinery and equipment supply their products with MTP files, and all common automation system manufacturers can both generate and import MTPs. Applications also exist outside the process industry, for example in shipbuilding. The technology is also suitable for integrating package units into non-modular plants, allowing centrifuges, blowers, temperature control units, and other equipment to be integrated seamlessly and cost-effectively.

5.3 Use Case: Semiconductor Industry

Semiconductor manufacturing is a special case of industrial production: over the course of manufacturing, a single wafer is divided into thousands of individual chips, while the required process steps are distributed across numerous countries and companies. Unlike in most industries, value creation therefore does not follow a traditional linear supply chain from one supplier to the next. Instead, it forms a highly fragmented, multi-stage lifecycle with many nodes and branches. The individual steps – from silicon and wafer supply through fabrication (FAB), sorting, die bank, assembly and testing, to integration, operation, and recycling – are distributed globally and regularly carried out by different companies. This network is coordinated by the chip designer, who acts as the orchestrating entity at the top and brings together the individual process steps (see Figure 3).

Semiconductor manufacturing depends not only on machines, but also to a large extent on chemical materials that, for example, make it possible to create the structures required for conductive paths. These materials are developed by the chemical industry together with semiconductor manufacturers and transferred into commercial production. They can have a significant impact on chip performance and production scrap rates. The influencing parameters cannot be analytically captured as clear cause-and-effect relationships. Data science and AI methods provide an opportunity to identify such relationships based on data. When variations or defects occur during production due to effects that cannot be clearly identified, their root cause cannot always be determined on the basis of these parameters alone.

A prerequisite for the trusted, data-driven identification of cause-and-effect relationships across company boundaries is therefore that the relevant data is available in an interoperable form – based on semantic descriptions that define a consistent data model that can be integrated by all participants. In Semiconductor-X, such data models are defined on the basis of the AAS, the Digital Reference, and SEMI.org standards. By incorporating existing standards into the common model, the highest possible level of interoperability across company boundaries is ensured.

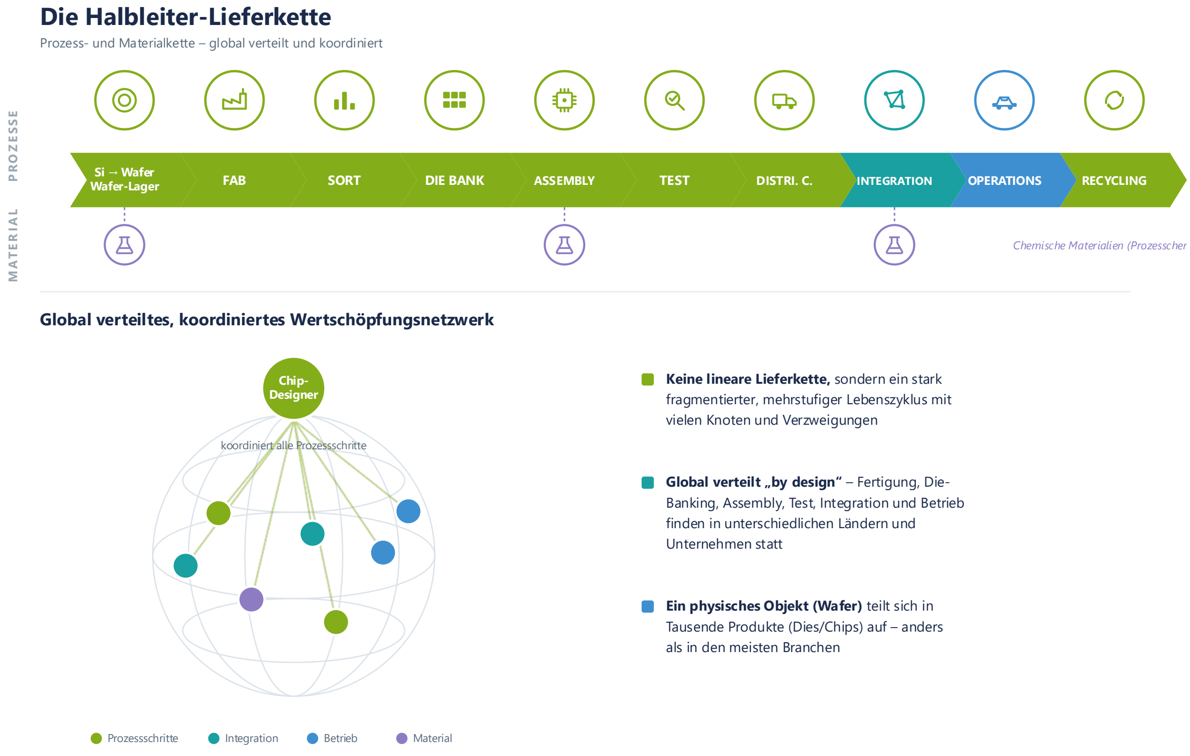

Figure 3: The Semiconductor Supply Chain

This becomes particularly relevant in the context of a fundamental development in chip design: the transition to chiplet-based designs. Whereas a complex chip was traditionally manufactured as a monolithic component on a single die, it is now increasingly divided into several smaller, functionally specialized components – chiplets – which can be manufactured at different technology nodes and in different fabs and subsequently integrated into a complete system via standardized die-to-die interfaces such as UCIe. This makes semiconductor manufacturing modular – but on two levels. At the product level, the chip itself becomes modular: it is composed of standardized, reusable chiplets that can be flexibly combined depending on the target product. At the production level, the fabs within the distributed manufacturing network themselves become modules – specialized, combinable manufacturing capacities that can be connected, exchanged, or reconfigured depending on the product, technology node, and capacity utilization. In the semiconductor industry, Modular Production therefore extends beyond the individual shop floor to become a cross-factory and cross-company network – a consistent extension of the modular paradigm beyond factory boundaries.

This potential becomes particularly evident with a new generation of AI chips, where the model or its weights are increasingly embedded directly into the hardware, for example in the form of dedicated accelerators or in-memory computing architectures. As AI model architectures evolve exceptionally quickly, monolithic chip designs with their long development cycles are coming under increasing pressure. A chiplet-based, modular approach, by contrast, allows compute, memory, and interface components to be recombined quickly and selectively. This is precisely where Modular Production contributes to the objectives outlined in the definition: it makes manufacturing faster, more scalable, and more flexible – characteristics that are becoming increasingly critical to success given the short innovation cycles of AI hardware.

To enable this distributed, modular interaction, Semiconductor-X relies on two key building blocks. The first is the **data space**, which enables sovereign and legally compliant data exchange between the participating companies. Using the Eclipse Data Space Connector (EDC), data can be securely shared between material suppliers, fabs, and integrators, embedded within access rights and usage management. This also addresses one of the central challenges: establishing in a legally compliant manner who owns the data and what each user is permitted to do with it. The technical demonstrator developed for Hannover Messe 2026 illustrates this approach: BASF, Merck, and a semiconductor fab exchange material, process, and product information via connected data spaces, with each company providing its data as a digital representation – for example, as an AAS of the chip.

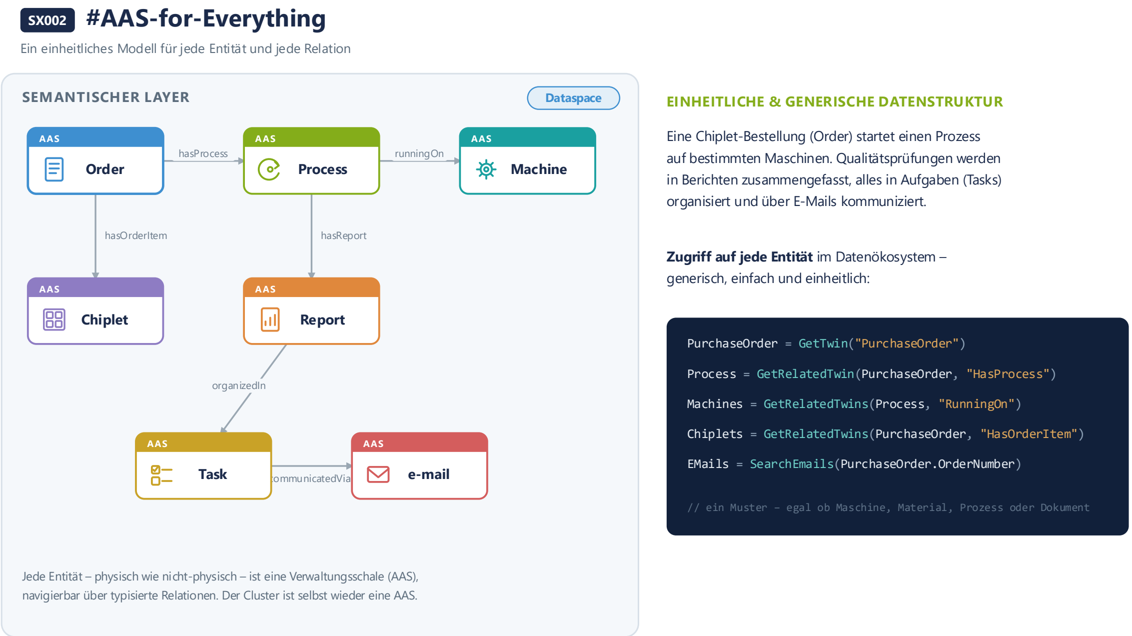

Figure 4: SX002 – #AAS-for-Everything

The second building block is the **#AAS-for-Everything** concept (see Figure 4). The basic idea is to use an Asset Administration Shell (AAS) not only to describe physical assets such as machines, products, and materials, but every relevant entity within the data ecosystem: orders, processes, reports, tasks, and even messages such as emails. Each of these entities is thus given a uniquely addressable digital twin with clearly typed relationships to other entities. Starting from a chiplet order, for example, the associated process, the machines executing it, the ordered chiplets, and the linked messages can all be identified in a consistent, generic manner. Access to the entire industrial data ecosystem therefore follows a single, simple, and consistent pattern – regardless of whether the entity is a machine, material, process, or document.

The **#AAS-for-Everything** principle results in an important structural characteristic. The information associated with a module is not contained within a single Asset Administration Shell, but forms a cluster of multiple individual AASs – one for each entity, for example the machine, the material being processed, the process being executed, the associated order, and the accompanying communication. This cluster is in turn represented as an Asset Administration Shell that bundles and references the individual shells and makes them externally accessible as a self-contained, modular data cluster. Asset Administration Shells can thus be nested within one another: a higher-level shell aggregates the shells of the entities it contains, which in turn may consist of further shells.

In this way, the modularity of production is directly reflected at the data level – just as physical modules consist of smaller components and can themselves be combined into larger units, their digital representations are also structured recursively. At the same time, this approach raises the question of how to manage the complex interconnection of Asset Administration Shells. How AASs are referenced, nested, and described in terms of their mutual relationships is a key design question that becomes increasingly important as the concept expands and has not yet been fully standardized.

These interconnected data clusters form a semantic layer – a digital graph model that links materials, processes, tools, products, and the supply chain. However, this semantic layer is far more than a "new generation" data structure. It is also inherently AI-friendly: because all entities and their relationships are described consistently and in vectorized form, AI can retrieve relevant data quickly and selectively instead of having to search and interpret extensive, unstructured data lakes. The semantic layer is therefore context-window-friendly: AI navigates along typed relationships and vectors and loads only the semantically coherent context actually required into its language model. This not only reduces ambiguity and the risk of hallucinations, but also makes large-scale operation significantly more cost-efficient, as computational and processing effort – and therefore token consumption – is limited to what is actually needed.

Agents can be deployed on top of these modular data clusters.[1] Each module is assigned its own agent, which operates on the module's data cluster, understands its capabilities and current state, and acts on behalf of the module. Because the agents communicate with one another via the data space and the common semantic description, they can coordinate production among themselves. Control is no longer handled centrally by a higher-level control instance, but in a decentralized, negotiation-based manner between the modules. Following this principle, capacity planning can take place directly at module level: through its agent, each module communicates available capacities, capabilities, and bottlenecks and independently coordinates the allocation of production orders with other modules. This concept extends the service, product, and resource agents described in the reference architecture to the cross-factory and cross-company level of semiconductor production, thereby building on the vision of digital twins and cooperating agents described at an early stage.[1] Based on the shared data foundation, the agents also enable traceable, AI-supported analysis of distributed data to identify clear cause-and-effect relationships between materials, processes, tools, and products – for example, to determine a "Golden Batch" or specifically reduce production scrap.

The benefits of this approach can be directly linked to the objectives of Semiconductor-X: Traceability, Zero Waste, and High Utilization. Consistent, semantically described data models provide the foundation for tracing every chiplet across the entire distributed supply chain, minimizing scrap through more precise, data-driven process control, and maximizing the utilization of distributed capacities. Because modular plants can already be developed and built according to these data models during the development phase, the result is an overall flexible and resilient supply chain that can respond quickly to changing market conditions.

6 Conclusion and Outlook

The Topic Group defines the cross-sector objective of Modular Production as follows: Modular Production describes a production approach in which plants, machines, and processes are built from standardized, reusable modules. Depending on the product, capacity requirements, or process changes, these modules can be flexibly combined, exchanged, or reconfigured. The objective is to break up rigid production structures and adapt production to new requirements more quickly, scalably, resiliently, and economically.

From a technical perspective, this is supported by standardized interfaces, digital twins, semantic descriptions such as the Asset Administration Shell (AAS), skill-based automation, and increasingly agent-based control.

All participating projects, from the process industry to discrete manufacturing, pursue this overarching objective through their respective approaches: increasing flexibility in production. On this basis, the reference model contributed by SmartFactory-KL was expanded to incorporate the perspectives of different industries, use cases, and technological focus areas.

Throughout the work of the Topic Group, specific industrial applications were examined. While Artificial Intelligence, particularly since the emergence of powerful Large Language Models, is increasingly finding its way into production, industrial self-control has so far been implemented only in selected applications. For the frequently described vision of a product "that manufactures itself," only a limited number of industrial approaches could be identified during the lifetime of the Topic Group, including applications in heating systems manufacturing, the audio and electronics industry, commercial vehicle production, and the kitchen furniture industry. In the process industry, the Module Type Package (MTP) is also a key topic for the future with significant potential. However, with the current trend toward Agentic AI and, in particular, Agentic Manufacturing, the industrial relevance of self-controlling production systems is expected to increase significantly in the coming years.

One industry has already been implementing the principle of self-control on a large scale for many years: the kitchen industry. Large production plants manufacture several hundred to more than a thousand kitchens per day, with each kitchen being a unique, customer-specific product. This is enabled by a consistently applied modular building-block principle that supports a high degree of product variety and numerous customer-specific adaptations, such as shortened cabinets or customized equipment. Production environments with very high product variety, a high degree of automation, and end-to-end digital planning are particularly well suited to self-control approaches. A key advantage of the kitchen industry is that the planning process already begins digitally with the end customer in the kitchen showroom. This early digital description of the product enables automated and autonomous production processes in the kitchen industry.

In addition, there are numerous examples in research, standardization, and technology transfer organizations that address self-control and Modular Production:

Organization / Project	**Organization / Project**	**Link**
Plattform Industrie 4.0, Working Group 3	White Paper *Collaborative Condition Monitoring*; *"On the Way to the Digital Champion 2030"*	LINK
OI4 Open Industry 4.0 Alliance	Series of white papers on *"Production of Tomorrow"*: self-control using digital twins – Arvatare	LINK

ARENA 2036	VWS4LS Project, Use Case 4: “Automation of Flexible and Modular Production Processes”	LINK
SmartFactory-KL	Demo factory & Whitepaper „Production Level 4“	LINK
ZVEI Zentralverband der Elektronischen Industrie	Whitepaper „Vom Automatisieren zum Agieren“	LINK

It is important to recognize that self-controlling production is not an isolated production topic. It begins as early as the development of products and production systems. Products must be described in a way that allows their requirements, variants, and required capabilities to be processed fully automatically in production. In the kitchen industry example, this is managed through complex master data. Capability- and skill-based production approaches play a central role here, as they link the configuration of customer-specific products with the available capabilities of plants and machines. In the context of Semiconductor-X, it also became clear that the digital material properties of a raw material can have a direct impact on production control. This extends the scope of Modular Production beyond the factory boundary to include materials, products, processes, plants, and the supply chain.

With regard to technological implementation, the participating projects agreed that the Asset Administration Shell provides a suitable format for this purpose. In particular, Type 3 of the Asset Administration Shell, as described at the beginning of this document, provides an appropriate foundation for describing and using industrial assets in connected production systems. At the same time, standardization gaps remain, particularly in the description of rules, events, and role models. To address these gaps, the Topic Group, together with other members of the Open Industry 4.0 Alliance, contributed to the development of a standard submodel as a proposal for the IDTA. If you are interested in the results of this standard submodel, please contact one of the four authors:

- Christian Methe (christian.methe@twindustrial.de)
- Haris Avgoustinos (haris@expo21xx.com)
- Pascal Rübel (pascal.ruebel@smartfactory.de)
- Thomas Tauchnitz (thomas@tautomation.de)

Looking ahead, the industrial use of Modular Production is currently visible only in selected applications, while the underlying concepts have already been widely developed and tested in research projects, demonstrators, and transfer initiatives. Agentic AI and Agentic Manufacturing are now providing a technological impetus that could significantly accelerate industrial implementation. For Europe in particular, which faces considerable pressure to innovate and improve efficiency in global industrial competition, Modular Production offers an opportunity to sustainably strengthen the flexibility, resilience, and competitiveness of industrial production.

At the time of writing this white paper, Plattform Industrie 4.0 is being strategically realigned toward Industrial AI and expanded with a new working group focusing on production. In addition, three of the participating companies intend to continue the topic of Modular Production both within this environment and within the Tractus-X community (LINK), with the aim of sustaining and further developing the work initiated by the Topic Group. If you are interested in contributing to this work or continuing the results, please contact the core team of authors.