Fraunhofer IOSB

visIT

Agentic Industrial AI

Concepts, applications and data engineering

www.iosb.fraunhofer.de

ISSN 1616-8240

Imprint

visIT is published about three times a year and informs about selected research topics of Fraunhofer IOSB. To order single issues, for a free visIT subscription, and for address changes and cancellations, please send an email to publikationen@iosb.fraunhofer.de

Publisher
Prof. Dr.-Ing. habil.
Jürgen Beyerer

Editor
Dipl.-Phys. Ulrich Pontes
Lena Kaul, M.A.

Layout
Anja Wollfarth, M.A.

Printing
Stober Medien
Industriestraße 12
76344 Eggenstein

Editorial address
Fraunhofer Institute of Optronics, System Technologies and Image Exploitation IOSB
Fraunhoferstr. 1, 76131 Karlsruhe
Phone +49 721 6091-300
presse@iosb.fraunhofer.de

Institute of the
Fraunhofer-Gesellschaft,
Munich, GERMANY

27th year
ISSN 1616-8240

Image sources
Cover picture: AI-generated visualization.

All other images
© Fraunhofer IOSB,
Exceptions are marked.

Topics

Editorial

Dear friends of Fraunhofer IOSB,

Following the rapid technological evolution in the field of Artificial Intelligence (AI), the title of this brochure – Agentic Industrial AI – seems quite straightforward. How can we apply the concepts and tools of AI agents to industrial applications? Indeed, there is a huge potential if optimization tasks in production processes can be delegated to autonomous software components, or if operators can reliably interact with machines or entire production plants in natural language. However, industrial environments have specific requirements, not only with respect to safety, security and reliability constraints, but also regarding the methods to acquire and manage the data needed to train and operate the AI agents, as well as to systematically engineer industrial AI systems. Hence, this brochure takes a broader perspective and includes methods, tools and applications that demonstrate how to integrate agentic AI into industrial IT environments based on Industry 4.0 standards and industrial dataspaces.

It starts with an essay about the revival of agent technology in the AI era, before presenting, in a guest contribution from EDI AG, an industrial AI framework product that is evolving towards agentic AI. We then consider the data foundations enabling agentic AI, based on open standards for sensor data management, Industry 4.0 digital-twin modelling standards, industrial dataspaces and data ecosystems. We continue with the question of how to apply large language models (LLMs) to industrial engineering before presenting one of the key interfaces for agentic AI – the Model Context Protocol (MCP). The brochure concludes with an explanation of how to use agentic decision support for power system control rooms.

Agentic industrial AI is highly dynamic, multidisciplinary and diverse. It requires the targeted and careful combination of systems engineering, AI and IT skills within and across companies. We would appreciate the opportunity to get in touch and stay in contact with you in order to collaborate on dependable and trustworthy agentic industrial solutions.

We hope you find this issue both informative and enjoyable.

Fraunhofer IOSB, April 2026

Prof. Dr.-Ing. habil. Jürgen Beyerer

Dr.-Ing. Thomas Usländer

Prof. Dr.-Ing. habil. Jürgen Beyerer

Dr.-Ing. Thomas Usländer

Revival of the agents: what is new in the AI era?

Dr.-Ing.
Thomas Usländer

Business developer AI
Systems Engineering

Phone +49 7243 992-480
thomas.uslaender
@iosb.fraunhofer.de

Flashback

Agents? Wasn't this a great idea 20–30 years ago, which then fell into oblivion? The fascinating promise behind it was the idea that a personal assistant, realized as a software component, could act on behalf of a human, just by delegating a task to it. Already in those days, it was obvious that such an "agent" needs a certain degree of "intelligence" and "autonomy". It should be sufficient to specify and request a task to be performed, accompanied by well-defined objectives. The exact steps for how to perform the task, even in case of problems or unexpected events, may be omitted.

But why did agents fail in those times? The simple answer: It was too early! The whole IT environment, from the performance of the underlying networking infrastructure up to the enterprise applications, was not yet mature and powerful enough for wide implementation and acceptance of agents. Artificial Intelligence (AI) systems were mostly symbolic, brittle and slow. The result: Although agent-based systems were complex, they were not "smart" enough to justify their complexity. In practice, simpler solutions, such as emerging web applications, did the job well enough. And, without a killer application, companies saw little reason to invest.

First, the software engineering community had to do their duties, e.g., modularization of distributed enterprise applications into (web and micro-)services, strategic and scalable use of the Internet resources (cloud computing), realization of the Internet of Things to learn about the environment and the "world" from deployed sensors of all kinds as well as countermeasures against the emerging cybersecurity problem. And, not to forget, the standardization of communication and middleware technologies as prerequisite to cope with the interoperability problem.

Today's AI agents

With today's AI capabilities, agents are experiencing a remarkable revival—and this time, the vision has a better chance to be successful if we carefully apply the AI agent technology.

One of the key drivers of this revival is the rise of large language models (LLMs). These systems can, or at least seem to, understand and generate human language, reason across domains and adapt to diverse tasks. Unlike earlier rule-based agents, these AI agents are not confined to predefined scripts. They try to interpret ambiguous instructions, ask clarifying questions and generate context-aware responses. This flexibility makes them far more capable and useful in real-world scenarios. This is, at least, the expectation.

Towards agentic AI

Another major innovation that is claimed by agentic AI is tool integration. AI agents today can connect to external applications, databases and APIs. Instead of merely producing text, they can perform actions: analyze reports and production logs, write and execute code or manage devices. Quick side note: This progress in tool integration largely benefits from the standardization work on communication protocols, services and related information models, e.g., the IEC 62541 OPC UA standard

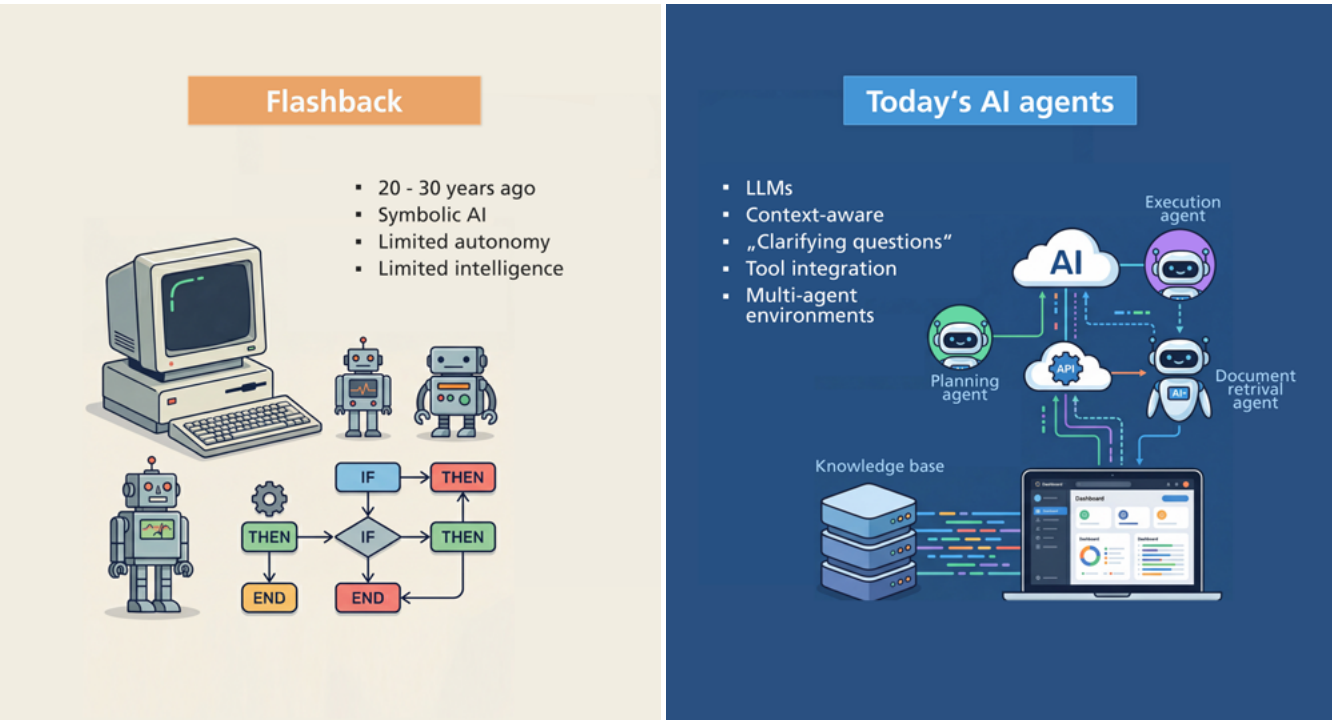

Fig. 2: Towards agentic AI – an infographic illustrating the evolution of artificial intelligence. (AI-generated visualization)

or the Industry 4.0 Asset Administration Shell and related information models in the last decades. Before, orchestration of diverse tool capabilities has been very expensive and error-prone. Hence, multi-agent systems are gaining importance. Instead of relying on a single monolithic model, developers now design agentic AI environments in which multiple specialized agents cooperate: one focusing on data acquisition, another on planning, another on execution.

However, the semantic embedding models of LLMs, i.e., conversion of text into numerical vectors, cannot fully resolve the problem of semantic interoperability, even more if the underlying data and documents are not yet semantically aligned. "Language is the source of misunderstandings" (Antoine de Saint-Exupery) is a quote that is also true for the design of these multi-agent AI systems.

Ethical and legal constraints

The revival of agents in the AI era also raises critical questions. Issues of data privacy, transparency, accountability and bias remain central concerns. When agents act autonomously, who is responsible for their decisions? How can we ensure ethical alignment with human values?

It is important to state that, although not yet explicitly defined in international standards, AI agents are variants of AI systems as these allow for varying degrees of autonomy and for influencing environments. This is exactly what "AI agents" do, which is why they naturally fall under the term "AI systems." Legally, the risk-driven approach of the European AI Act, that is centered around the term "AI system", applies. Depending on the risk classification of an AI system, actions must be taken to eliminate or mitigate risks, with the aim of building trust and ensuring that AI agents benefit society as a whole.

In conclusion, the revival of agents in the AI era is not merely a technological upgrade—rather, it represents a paradigm shift. AI agents are evolving from simple assistants into intelligent partners. With a systematic AI systems engineering approach, Fraunhofer IOSB does contribute to a demand-driven and sustainable use of AI technologies in complex technical environments [1]. This includes the careful design of cooperative industrial AI agents, where necessary for solving the problem.

1 Usländer, T. (2025). KI-Engineering in industriellen Datenräumen. In: Hoffmann, C.H., Hersberger, S. (eds) Wie die Künstliche Intelligenz die Wirtschaft verändert. Springer, Wiesbaden. https://doi.org/10.1007/978-3-658-46839-2_2

Agentic industrial AI for cross-layer automation

Unlocking data and knowledge silos with the EDI hive IoT & AI Framework

Dr.-Ing. Mohanad El-Haji

EDI AKTIENGESELLSCHAFT

Phone +49 721 79199-155
el-haji@edi-ag.ai
www.edi-ag.ai

Dr.-Ing.
Thomas Freudenmann

EDI AKTIENGESELLSCHAFT

Phone +49 721 79199-155
freudenmann@edi-ag.ai
www.edi-ag.ai

In many industrial organizations, valuable operational data and expert knowledge remains isolated in systems, departments or individual expertise, limiting transparency and insight reuse.

Agentic industrial AI breaks these barriers by accessing heterogeneous data sources, interpreting them within domain-specific contexts, and autonomously supporting or executing operational decisions. This unifies automation across all layers of the industrial automation pyramid, linking enterprise systems, e.g., ERP and MES, to sensor-based shopfloor infrastructure.

The EDI hive IoT & AI Framework

To power such systems, EDI created the EDI hive IoT & AI Framework, a modular platform for agentic industrial AI applications. Developed over a decade, the framework combines IoT technologies, semantic data spaces and modern AI methods such as Retrieval-Augmented Generation (RAG). Its architecture aids:

- real-time industrial data streaming
- automated semantic integration of varied data sources
- AI model and agent deployment across IT and operational technology (OT) environments
- secure user and access management for industrial data

The framework bridges industrial data and knowledge, from enterprise applications to shopfloor processes, ensuring consistent cross-layer automation.

The EDI hive IoT & AI Framework introduces three core innovations that enable scalable agentic industrial AI:

- **cross-layer semantic integration**
  EDI hive joins diverse industrial systems

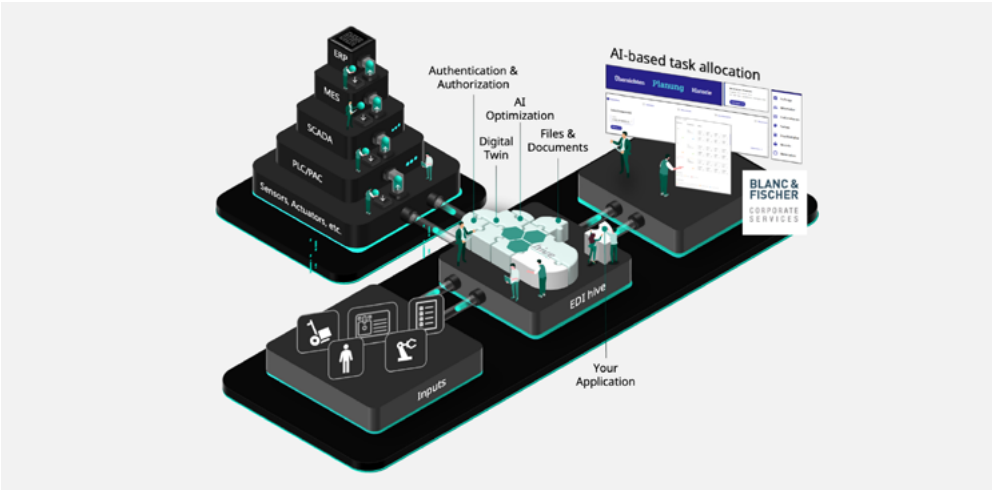

Fig. 1: AI agents for seamless automation across the layers with EDI hive IoT & AI Framework (Source: EDI AG).

Fig. 2: AI agent–supported service workflow with domain-specific data translation (Source: EDI AG).

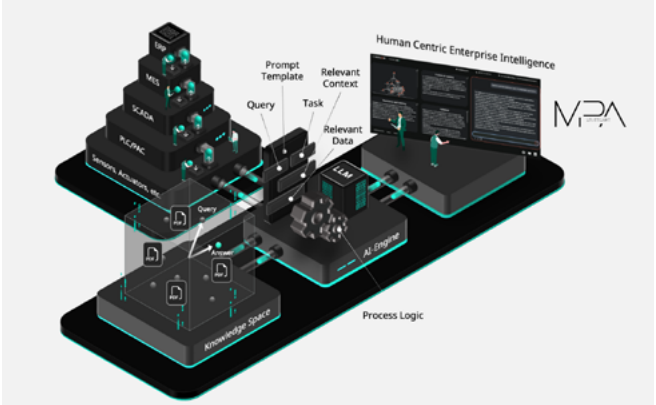

Fig. 3: AI agent supporting process development and production of component connections (Source: EDI AG).

across all layers of the automation pyramid. Beyond data aggregation, the platform interprets and contextualizes system information, allowing AI agents to orchestrate workflows from enterprise IT to OT.

- **causal knowledge representation through cause-and-effect chains**
  Cause-and-effect chains, a key patented innovation, formalize expert knowledge by mapping causal relationships among process parameters, operational conditions and system outcomes. In EDI hive, these chains form a meta-ontology for causal AI, encoding domain expertise in machine-readable formats. It allows AI systems to offer transparent reasoning and actionable industrial processes insights.
- **agentic AI with retrieval-augmented knowledge systems**
  EDI hive integrates RAG-based AI agents that retrieve relevant information from distributed knowledge bases. Using LLMs, these agents generate context-aware recommendations or decisions for industrial workflows.

Industrial applications

The EDI hive IoT & AI Framework has proven successful in multiple industrial applications, illustrating the potential of agentic AI for cross-layer automation.

- **workforce planning in manufacturing**
  In the KARL project [1], an AI agent assists shift supervisors in dynamic personnel scheduling. It retrieves SAP production planning data, optimizes workforce allocation and generates fair, efficient shift plans.
- **smart customer service**
  An AI agent analyzes machine log data from grinding systems and integrates it into a Microsoft Dynamics 365 ticketing system. This creates a seamless workflow from shopfloor equipment to service management, improving response times and service transparency.
- **engineering knowledge and process development**
  In projects such as CyberJoin [2] and AnAttAI [3], RAG-based knowledge systems assist engineers in complex manufacturing processes such as automotive aluminum welding. By combining experimental data, expert assessments and scientific publications, the AI agent provides ranked technical recommendations to support engineering decisions.
- **AI-driven digital twins for predictive process control**
  Cross-layer automation can be extended with digital twin and AI-based hybrid sensors. In waste-to-energy plants using the StableFlame [4] process, hybrid sensors convert image data into process-relevant KPIs that are integrated into a digital twin. Predictive models then optimize operations, e.g. crane automation, and continuously improve process.

Toward the next generation of industrial automation

Agentic industrial AI represents a new paradigm for intelligent automation. By integrating semantic system connectivity, causal knowledge representation and AI agents, the EDI hive IoT & AI Framework supports organizations to unlock distributed expertise and deploy cross-layer automation in complex industrial environments. It sets the technological stage for adaptive, knowledge-driven industrial systems of the future.

1 https://kompetenzzentrum-karl.de/ki-im-einsatz/einsatzplanung.
2 https://cyberjoin.de/angebot/.
3 https://www.materialdigital.de/project/23.
4 https://www.stableflame.com.

Unlocking AI potential

Open standards as the foundation for AI agents

From Chatbots to AI agents

Hardly any technology dominates the current discussion as much as Large Language Models (LLMs): AI models that understand and generate text. To unlock the real potential for AI, we need to go one step further to AI agents. While a Chatbot answers a question and the interaction ends there, AI agents follow a fundamentally different cycle (cf. Figure 1): They **perceive** by ingesting data from their environment. They **decide** by autonomously planning which steps are necessary. They **act** by actively intervening in their environment or proposing actions. And they **iterate** by evaluating intermediate results and adjusting their plans. At every step of this cycle, the AI agent interacts with its surroundings. As soon as it is expected to create value in the real world, it must be able to clearly understand and interpret that environment. This requires that the data describing its environment is provided reliably, in a structured
manner and in a machine-interpretable form. Without this foundation, the AI agent is blind.

FAIR principles and their limits for AI agents

The FAIR principles have become the internationally recognized guiding framework for data provision. FAIR encapsulates four tenets of good data management: data should be **Findable** through rich metadata and persistent identifiers, **Accessible** via standardized protocols, **Interoperable** through shared vocabularies and formats, and **Reusable** under clear licensing and provenance information. FAIR has sharpened the focus on metadata and accessibility, creating an indispensable foundation. However, FAIR was designed for

Philipp Hertweck

Information Management and Production Control (ILT)

Phone +49 721 6091-372
philipp.hertweck
@iosb.fraunhofer.de

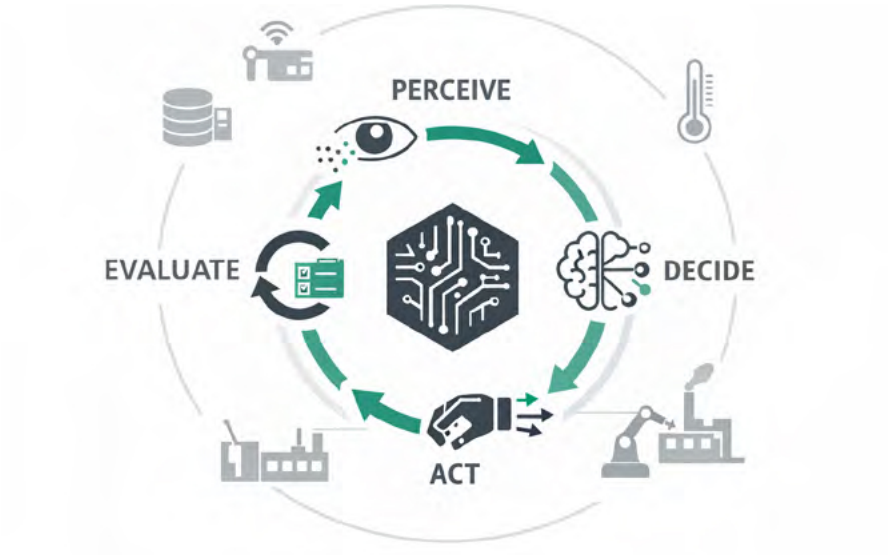

Fig. 1: AI agents operate in a continuous cycle of perception, decision-making, action, and evaluation.

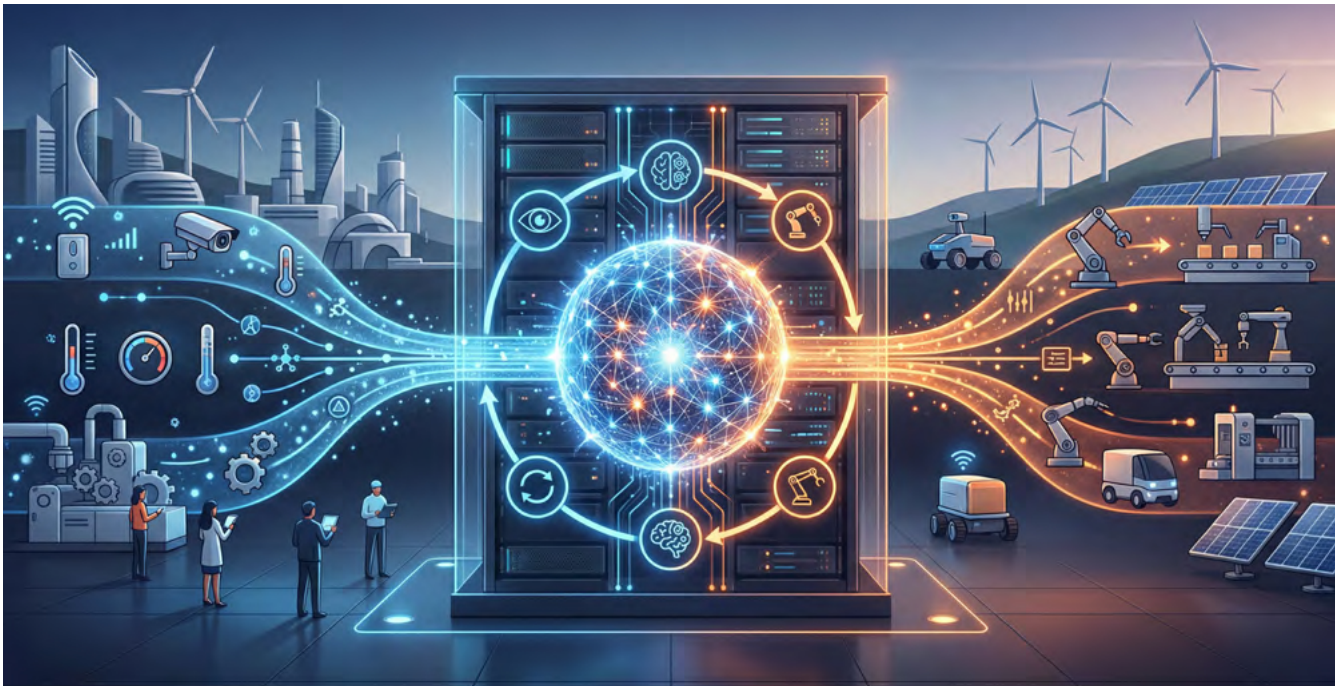

Fig. 2: Open standards are the foundation that enables AI agents to understand and interact with the real world. (AI-generated visualization)

human users who search for, evaluate and manually prepare datasets. In a world of autonomous AI agents, the focus shifts: It is no longer the description of the data that determines its usability, but the data itself—its structure, its embedded semantics and the way it can be consumed via interfaces. FAIR is an indispensable precursor. But FAIR alone does not make data AI-ready.

AI-ready means: An autonomous software actor can query, interpret and use the data for decisions without human intervention. This is exactly what an AI agent needs to make its action cycle work. Between FAIR and AI-ready, concrete gaps exist: AI-ready requires **standardized APIs** rather than merely open protocols, **consistent data models** rather than just metadata schemas, **embedded semantics** within the data model rather than externally documented vocabularies—and the **ability to interact** with the real world, which FAIR does not address at all.

How open standards bridge the gap to AI-ready data

This is precisely where open standards come into play, not as a single technology, but as a design principle. In practice, there will not be just one data source and one agent. Numerous sources and AI agents must interoperate to respond flexibly to new requirements. This can only succeed if data provision and the AI agent agree on common interfaces. Open standards prevent silos by defining vendor-neutral interfaces. They ensure that data from different sources can be directly compared and combined. They normalize data models, anchor semantics within the standard itself and ideally enable not only data retrieval but also the standardization of actions.

One example: The **OGC SensorThings API** (consisting of two parts: Sensing and Tasking), an open standard for IoT sensor data management, addresses all the gaps mentioned above. It offers RESTful access with JSON serialization and OData queries, defines a consistent data model from Things through Datastreams to Observations, anchors semantics through defined ObservedProperties and, with Part 2 (Tasking), enables the control of IoT devices. The agent thus evolves from a passive data consumer into an actionable actor.

Ultimately, an AI agent can only be as good as the context provided. This context describes access to data that reflects the real world. The future of AI agents lies not only in better models, but in a better understanding of the environment—and that understanding begins with open standards.

Data ecosystems as enablers of industrial AI

Dr.-Ing. Olaf Sauer

Business Developer
Automation and
Digitalization

Phone +49 7243 992-477
olaf.sauer
@iosb.fraunhofer.de

In addition to established world-class hardware-related competencies, factory operators and equipment suppliers must develop capabilities to effectively apply artificial intelligence (AI), machine learning (ML), digital twins and data ecosystems. Sustainable value creation through AI cannot be achieved in isolation but requires cooperation across the entire value chain, with partners contributing complementary expertise.

Data ecosystems are capable of sharing data securely from various sources such as machine data, quality data or maintenance data and provide easy access via standardized connectors. This enables the development and reuse of AI applications across organizational and system boundaries. Standardized formats and metadata contribute to data accessibility and quality and enable the targeted use of data from the supply chain or from operations. Open standards facilitate the integration of data from diverse systems such as ERP, MES, SCADA or PLM.

Companies are often reluctant to share data across organizational boundaries due to concerns about disclosing proprietary knowledge. Data ecosystems address these concerns through clearly defined governance and data security mechanisms, including identification and authorization, access and usage control, trusted data sources and automatable data access agreements.

As a result, ecosystem partners such as suppliers, customers and service providers are able to integrate data from entire value chains and exploit benefits they would never be able to leverage on their own.

Benefits of data ecosystems for industrial AI

- Prepared data sets can be reused across multiple AI applications, including predictive maintenance, process optimization and quality prediction.
- AI models achieve higher quality and trustworthiness, as consistent data improve model accuracy and contribute to the stability and comprehensibility of AI applications.
- Sharing data within trusted data ecosystems avoids isolated solutions and enables the gradual, risk-reduced deployment of AI, from pilot projects to serial production.
- Open standards reduce data preparation effort per project, enabling faster prototyping and iterative optimization of AI applications.
- Partnership-based value creation emerges when companies along a value chain use data cooperatively, for example by jointly training models with suppliers, thereby strengthening the competitiveness of all participants in the data ecosystem.

Typical application scenarios

- Predictive maintenance, where AI identifies impending machine or component failures early, enabling optimized maintenance scheduling and reduced downtime.
- Process optimization, for example through the fine-tuning of process parameters to improve the learning curve of new processes.
- Quality predictions, enabling early detection of deviations and the reduction of scrap and rework.
- Energy and resource efficiency, with AI

Fig. 2: Digital solutions for real-world problems enable factory managers to take better and faster decisions. (AI-generated visualization)

models identifying inefficiencies and improving consumption profiles.
- Digital twins of individual plants or lines, combining runtime data from machines and systems with data from physical simulations to enrich models and generate additional training data for machine learning.

As a result, AI applications can be implemented faster, at lower cost with reduced risk, leading to improved availability, quality and production efficiency and ultimately improved competitiveness.

Data Space Lab act as AI incubators

Many factory operators and equipment suppliers are small and medium-sized enterprises with limited IT resources, limited space for prototyping and almost no free capacity for data acquisition from running machines. Consequently, the development and deployment of AI and ML in data ecosystems typically requires collaboration with trusted partners. This promotes cooperation and innovation within the ecosystem and paves the way for new business models that were previously unprofitable.

This is precisely why Fraunhofer IOSB is preparing its established research factories and AI real-world laboratories in Karlsruhe [1] and Lemgo [2] to serve as Data Space Lab („Datenraumwerkstätten" in German): Here, manufacturing companies will find everything they need to quickly and easily develop and test data space technologies, and the added value provided by the data ecosystem becomes clear and understandable.

Elements a Data Space Lab should offer:

- Options for instrumenting machines, systems and production processes where existing or planned sensor technology is insufficient for reliable data acquisition.
- Processing of manufacturing data in the edge-cloud continuum, combining on-site computing resources with a secure and scalable cloud infrastructure, e.g. to train and create AI- and ML-models for predicting quality, availability or output of components and machines.
- MX-Port configuration tools for rapidly connecting assets to the data space via the different MX-Port implementations [3] (currently Hercules—EDC, Leo—Asset Administration Shell (AAS) and Orion—OPC UA).
- Expertise in IT security management and IT security, including IEC 62443, network and communication security, IT-/OT-integration, and access and usage control mechanisms.
- Systematic development and operation of AI applications in industrial data spaces [4], following the established AI engineering methodology [5].

Based on these elements, Data Space Lab provide different collaboration formats, including bilateral support for individual companies, hackathons for the collaborative testing of connectors and governance workshops in which participants jointly define rules, roles and responsibilities for a data space.

---

1 https://www.reallabore-innovationsportal.de/karlsruher-forschungsfabrik-fuer-ki-integrierte-produktion, last access February 16, 2026.

2 https://www.smartfactory-owl.de/en/about-us/, last access February 16, 2026.

3 https://factory-x.org/wp-content/uploads/MX-Port-Concept-V1.10.pdf, last access January 26, 2026.

4 Usländer, T.: KI-Engineering in industriellen Datenräumen. In: Hoffmann, C. H., Hersber-ger, S. (eds.): Wie die Künstliche Intelligenz die Wirtschaft verändert. Springer: Wiesba-den, 2025. https://doi.org/10.1007/978-3-658-46839-2_2.

5 Usländer, T. and Schulz, D. (eds.): KI-Engineering in der Produktion. Whitepaper der Fraunhofer-Institute IOSB und IAIS. Fraunhofer Verlag: Stuttgart, 2023. https://doi.org/10.24406/publica-1685.

Agentic AI in the FA³ST Ecosystem

Marc Leon Haller

Information Management and Production Control (ILT)

Phone +49 721 6091-672
marc.haller
@iosb.fraunhofer.de

Digital twins have become a cornerstone of the digital transformation of German industry, as they enable digital representations that integrate assets across the industrial value chain into digital processes.

From the perspective of industry stakeholders, digital twins are no longer a topic for the future. According to a recent survey by Bitkom, nearly 50 percent of industrial companies already apply digital twins. Moreover, a clear majority of the surveyed companies consider them essential for maintaining competitiveness in international markets [1].

This assessment is hardly surprising. Digital twins form the foundation for digital business models such as Manufacturing-as-a-Service and support the efficient implementation of regulatory requirements, for instance in the context of Digital Product Passports.

Such cross-organizational applications require standardized information exchange across heterogeneous systems. In this context, the Asset Administration Shell (AAS) has emerged as a standard for interoperable digital twins to meet this requirement. The AAS defines a standardized metamodel for describing assets, specifies uniform interfaces for accessing and exchanging data and enables the provision of semantic information for consistent interpretation among different systems.

The FA³ST Ecosystem

While the AAS provides the conceptual foundation for interoperable digital twins, its practical adoption remains challenging. Companies often face the complexity of the specifications, the high implementation effort and the identification of relevant information from heterogeneous data sources required for modelling AAS-compliant digital twins.

To address these challenges, the FA³ST Ecosystem provides a modular and open set of tools that support companies of different sizes in implementing AAS-based digital twins [2]. As shown in Figure 1, the ecosystem comprises several artefacts that can be flexibly combined depending on specific requirements and use cases. For example, the FA³ST CreAitor supports the automated modelling of AAS-based digital twins, while the FA³ST Service enables their deployment and interaction [3].

Applying AI agents in the FA³ST Ecosystem

AI agents are integrated into the FA³ST Ecosystem to simplify complex and time-consuming tasks. For example, they can be applied within the FA³ST CreAitor to support the automated modelling of AAS-based digital twins. As illustrated in Figure 2, the agent assists during information processing by identifying relevant information from heterogeneous data sources such as technical datasheets, machine documentation or company websites. It analyzes, extracts and structures the input data, and links the extracted information to definitions provided by domain-specific ontologies.

Human-in-the-Loop

While AI agents can simplify complex and time-consuming tasks, an important limitation must be considered. Large Language Models

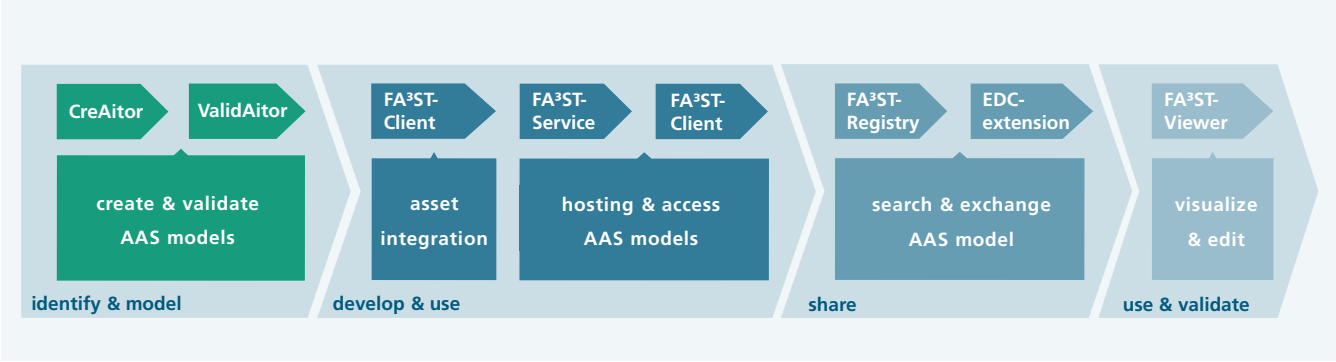

Fig. 1: Artefacts of the FA³ST Ecosystem.

underlying these agents are primarily trained on publicly available data and therefore lack domain- and company-specific knowledge.

As shown in Figure 2, the FA³ST CreAitor addresses this limitation by integrating additional knowledge bases into the agent's workflow. Domain-specific ontologies provide formal definitions and relationships that support the agent during information processing. In addition, users can be directly involved in the workflow to contribute contextual knowledge. Within the FA³ST CreAitor, users can specify their application domain, review extracted information or provide custom ontologies as additional knowledge.

This combination of formal domain knowledge, human expertise and AI agents ensures that AAS-based digital twins generated by the FA³ST CreAitor are consistent with domain requirements and company-specific conventions.

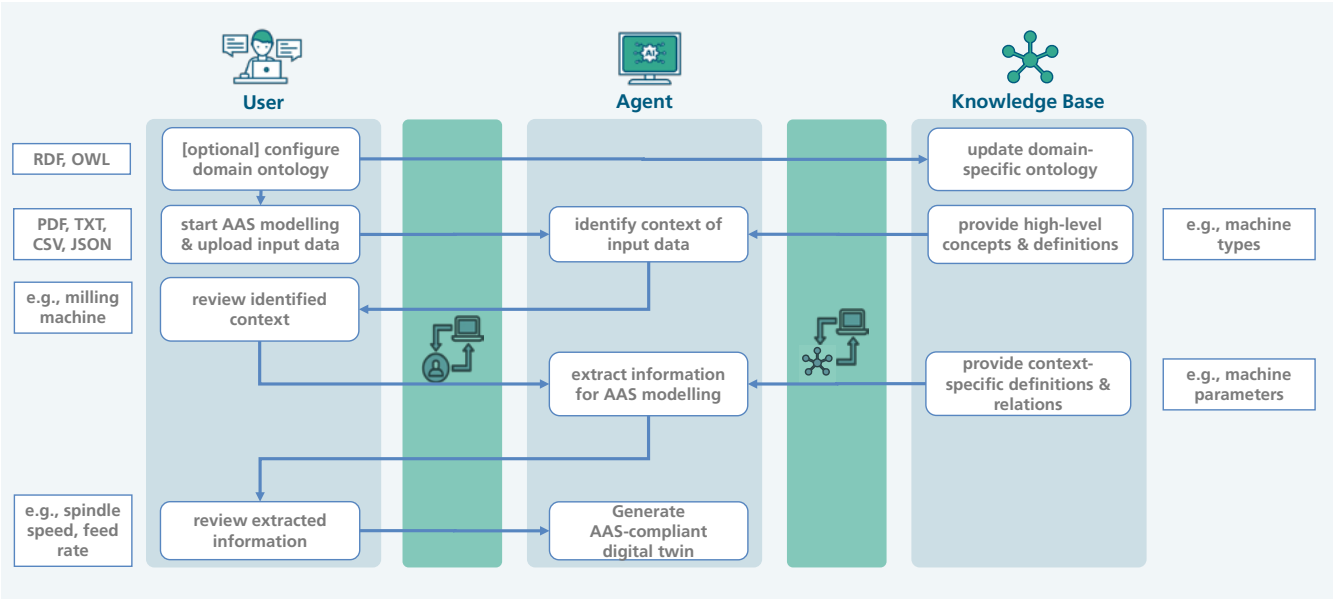

Fig. 2: Applying AI agents during automated modelling of AAS-based digital twins.

1 Bitkom: Industrie 4.0. Studienbericht. 2025. http://dx.doi.org/10.64022/2025-industrie-4-0.

2 More information about the FA³ST Ecosystem is available under: https://www.iosb.fraunhofer.de/faaast.

3 More information about the FA³ST Service is available under: https://github.com/FraunhoferIOSB/FAAAST-Service.

LLM agents for industrial engineering

From code generation to process optimization

Dr. Joschka Kersting

Machine Intelligence (MIT)

Phone +49 5261 9429-96
joschka.kersting
@iosb-ina.fraunhofer.de

Agentic AI describes systems that observe outcomes, evaluate success and adjust their approach autonomously. In industrial engineering, such systems are increasingly being tested for practical applications—also at Fraunhofer IOSB, where we develop and evaluate agentic solutions for industrial automation and manufacturing contexts.

How AI agents can support engineering work

The appeal of agentic systems lies in their ability to handle tasks that previously required engineers to navigate fragmented information manually. When a maintenance technician needs to understand why a machine behaves unexpectedly, the answer often lies scattered across maintenance logs, vendor documentation and historical incident reports. An AI agent can search these sources, correlate findings and suggest probable causes—not replacing human judgment, but accelerating the diagnostic process. Earlier AI applications were often simple chatbots, unable to help navigate extensive data sources or complex task histories. In contrast, agentic systems now provide targeted analysis and processing of industrial tasks, delivering real value to engineering.

Similar benefits emerge in quality management, where agents compare production parameters against specification limits and flag deviations. Agents can also help create structured documents—such as shift summaries, compliance documentation or audit preparation—from unstructured information. At Fraunhofer IOSB, we are exploring how such capabilities can be tailored to specific industrial domains.

Code generation for industrial controllers

Programming PLCs (Programmable Logic Controllers) remains a specialized discipline. Each vendor implements proprietary dialects of IEC 61131-3 Structured Text with incompatible function blocks and syntax conventions. Standard LLMs, trained on general-purpose programming languages, lack this domain knowledge.

Current approaches use compiler-in-the-loop architectures: The agent generates code, submits it for compilation, analyzes error messages and refines the output iteratively (cf. fig. 2). Fine-tuned models running on standard hardware achieve compilation success rates well above 85 percent. Moreover, agentic AI helps with the maintenance of legacy code bases in languages less known to today's developers (e.g., AWL/IL).

Formalizing optimization problems through dialogue

Manufacturing involves many optimization tasks: production scheduling, resource allocation, logistics planning. Mathematical solvers exist, but formalizing problems in solver-compatible form requires expertise that many companies lack.

One approach uses LLM-based agents to conduct structured dialogues with production planners. The agent asks targeted questions, identifies missing constraints and constructs a formal model step by step. The quality of the resulting models depends heavily on how clearly users articulate their

Fig. 1: Agentic AI in practice: aiding industrial automation with code generation at the engineering workstation; Fig. 2: Output of an agentic coding web app with ST code and compiler result (German).

requirements—implicit knowledge and edge cases are difficult to capture automatically.

Retrieval-Augmented Generation for technical knowledge

Companies accumulate documentation over years: maintenance logs, process specifications, compliance records. Retrieval-Augmented Generation (RAG) enables agents to access this knowledge and ground their responses in proprietary data rather than generic training information.

Challenges include retrieval quality for technical documents, which is a majore enabeler for a projects success. Our institute tunes dedicated systems to get most of small data collections for retrieval quality. Challenges also include handling multimodal content such as circuit diagrams or piping schematics. When retrieval returns irrelevant passages, the agent may generate plausible-sounding but incorrect answers—a failure mode that requires careful monitoring in practice.

Outlook

Agentic AI for industrial applications is moving from prototypes toward initial deployments. Integration with existing automation infrastructure, validation in safety-critical environments and robustness under real-world conditions are central challenges. Fraunhofer IOSB takes on these challenges every day and helps firms get actual benefits from recent technologies. Practical experience over the coming years will show where these systems deliver value and where their limitations lie.

MCP as an open interface for intelligent production

Jan Hermes

Cognitive Industrial Systems (KIS)

Phone +49 721 6091-435
jan.hermes
@iosb.fraunhofer.de

Benedikt Stratmann

Cognitive Industrial Systems (KIS)

Phone +49 721 6091-428
benedikt.stratmann
@iosb.fraunhofer.de

The Teuken4OPC project shows how modern AI assistants can be securely and scalably connected to existing production systems without proprietary solutions or deep changes to control systems. The Model Context Protocol (MCP) connects large language models (LLMs) with industrial standards such as OPC UA and Asset Administration Shell (AAS), and can also link to other systems like databases or sensor buses.

Initital situation: lots of data, little intelligent use

Many factories already have extensive data from OPC UA servers, AAS models, databases, historian systems and sensor buses. What is often missing is an intelligent layer that links this information and turns it into actionable, natural-language recommendations.

MCP as an open bridge

MCP is a protocol initiated by Anthropic and designed as an open standard, which is currently being further developed collaboratively by industry and open-source stakeholders. It provides LLMs with secure and traceable access to external tools and data sources via a vendor-independent interface.

In Teuken4OPC, MCP acts as a hub between AI and OT/IT: instead of connecting each machine directly to an LLM, MCP tools are provided once (e.g., an OPC UA tool for reading/writing variables and an AAS tool for accessing digital twins). The AI agent uses these tools to query sensor values, compare parameters or interpret process information, turning classic industrial interfaces into intelligent service interfaces—without modifying existing systems.

From OPC UA and AAS to databases and sensor buses

In the project, OPC UA and AAS were initially implemented as the first integration points. OPC UA is used to connect classic automation technology such as controllers, sensors and actuators, while AAS provides structured information and metadata about systems and products. On this basis, the AI agent is already capable of accessing process variables to answer natural language queries today, for example with the question "What is the current fill level in the tank?". It can also automatically identify relevant variables and sensors and propose step-by-step plans for analysis or to support operators. The architecture is open by design, so databases, historian systems or sensor buses can be connected via MCP later. In this way, previously isolated stand-alone solutions evolve into a scalable assistance system that unifies information across sources and provides it in an understandable form.

Example applications from the factory

Three very different scenarios were implemented in the project.

- **Process engineering:** A test plant with tanks, pump and valve is connected via OPC UA. The AI assistant can, controlled via natural language, read sensor data, explain states and perform simple control tasks.

Fig. 1: Vision of smart manufacturing: Machines deliver insights, performance forecasts and process analysis, with results explained in natural language by an LLM agent.

- **Injection molding:** A standard injection molding machine is connected via OPC UA. The assistant has read-only access and explains the relationships between process parameters and part quality. It indicates potential optimization areas and helps users understand the extensive range of machine variables without adjusting the machine itself.
- **Pick & Pack:** For a robotic system for automated packaging, objects are modeled as AAS. Via MCP tools, the AI assistant retrieves 3D geometries and metadata, suggests optimization factors, and generates packing plans for automated or manual execution.

Together, these use cases demonstrate that MCP serves as a universal bridge and unfolds its strengths particularly in combination with an agent-based AI assistant, whether in traditional process engineering, high-end injection molding systems or flexible robotics.

Scalable architecture instead of single-point solutions

A central outcome of Teuken4OPC is an IT architecture concept that goes beyond individual systems.

Latency and security critical functions run directly on the shop floor; knowledge- and computation-intensive components (LLMs, knowledge stores or monitoring systems) run centrally in the data center.

Combining agent-based AI logic with MCP as the connecting layer allows stepwise integration of new systems, databases or sensor buses without rebuilding the entire system. This creates a path from the first pilot application to a site-wide, scalable AI assistance.

Outlook: from demo to widespread application

Even today, the MCP-based solution facilitates access to complex systems: Users can ask questions, get process explanations, or receive decision support without being OPC UA or AI experts.

As LLMs evolve and more systems (databases, sensor buses, cloud services) are integrated, MCP can serve as a key technology for gradually enriching existing production environments with intelligence—standards-based, extensible and future-proof. At the same time, the agent-based system design ensures safe, transparent and effective operation.

Teuken4OPC thus creates a practical foundation for how today's interfaces can become the interfaces of tomorrow's intelligent assistants.

1 More info on: https://www.anthropic.com/news/model-context-protocol.

GridCompanion

Agentic decision support for power system control rooms

Power system control rooms, such as the grid control laboratory at Fraunhofer IOSB-AST (see Fig. 1), operate in an environment characterized by high system complexity, strict safety requirements and increasing time pressure. Operators must continuously interpret SCADA[1] measurements, topology changes, switching states and disturbance reports in order to maintain secure system operation. Although advanced analytical tools exist, they are typically used in isolation, require manual coordination and demand significant expertise to interpret under time-critical conditions.

From isolated tools to integrated decision support

GridCompanion [1] introduces an agentic decision-support architecture designed to address this fragmentation and to augment, rather than replace existing operational processes, see Fig. 2. The system does not autonomously intervene in grid operation. Instead, it acts as a goal-directed orchestration agent that coordinates deterministic analysis tools, executes them via standardized interfaces to existing control room systems and interprets their outputs in an operator-centered manner.

At its core, the agent is connected to the productive SCADA environment through a standardized interface that provides structured access to current grid states, measurements and topology information via the CIM[2] / CGMES[3] format. This ensures that all analyses are grounded in validated, real-time operational data. The agent itself does not replace established grid analysis engines; rather, it leverages them. Deterministic tools such as power flow calculation, contingency analysis or constraint evaluation modules are invoked through well-defined

Dr.-Ing. Dennis Rösch

Cognitive Energy Systems (KES)

Phone +49 3677 461-188
dennis.roesch
@iosb-ast.fraunhofer.de

Fig. 1: In the grid control lab of Fraunhofer IOSB-AST, the first practical tests with the GridCompanion are being performed.

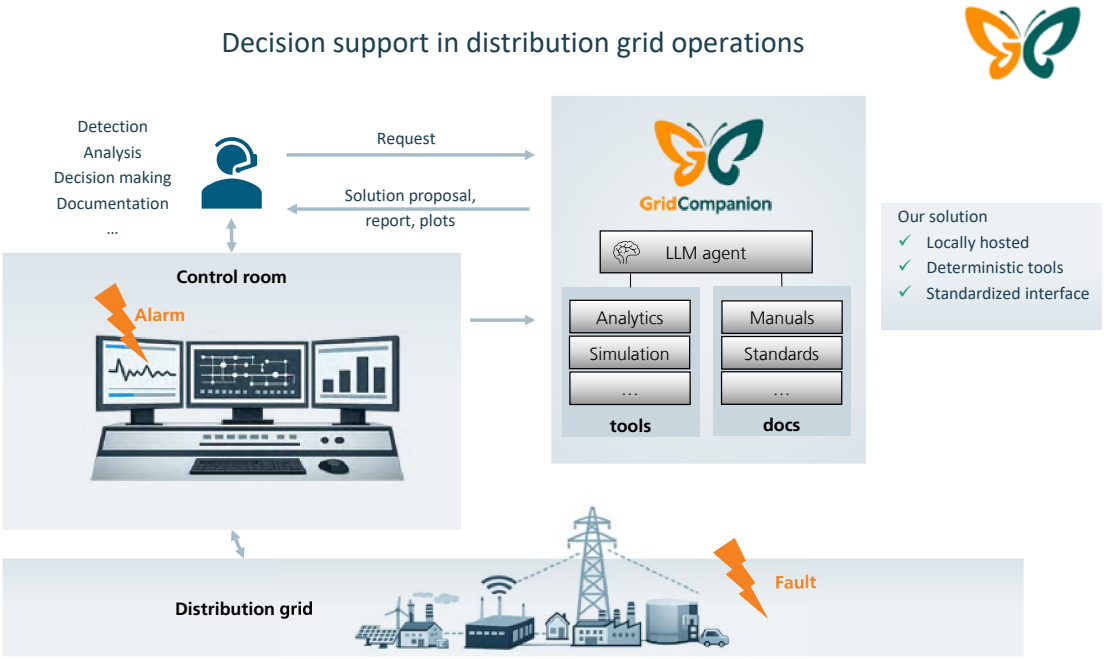

Fig. 2: Architecture of GridCompanion, an LLM-based decision support system supporting distribution grid operations.

APIs[4] (and later via MCP[5]) so that their numerical results remain reproducible, auditable and aligned with existing engineering practices.

An agentic orchestration layer for grid analysis

The agentic capability emerges from the orchestration layer. Given an operational objective, for example, assessing the impact of a disturbance, the agent selects appropriate deterministic tools, defines execution sequences and aggregates intermediate results into meaningful, task-specific views. A locally deployed LLM[6] provides the reasoning layer that structures this workflow: It translates high-level operator intents into tool calls, manages parameterization, maintains contextual state across analysis steps and synthesizes outputs into coherent, human-readable explanations and recommendations.

In addition to operational analytics, GridCompanion integrates a research agent connected to a RAG[7] pipeline and a knowledge base (e.g., operational guidelines, protection and control documentation, incident reports). Retrieved documents are combined with LLM-based reasoning to provide context-aware, sourced answers to operator questions. All evidence remains inspectable, and references to underlying documents are made explicit. Transparency and traceability are central design principles. Every recommendation generated by the agent is linked to its underlying deterministic results and, where applicable, to supporting documents. The complete reasoning chain, including tool invocations, parameters, intermediate outputs and cited knowledge sources, can be inspected step by step. The agent does not perform switching operations, modify setpoints, or trigger automated control actions. All final decisions remain with human operators in a human-in-the-loop manner.

This architecture reflects a shift from static dashboards and isolated tools toward dynamic, agentic orchestration. Instead of simply presenting raw outputs from individual applications, GridCompanion contextualizes findings across tools and knowledge sources, highlights constraint violations, compares alternative scenarios and formulates prioritized decision options tailored to the operator's current task, thereby supporting faster, more informed decision-making in the control room.

List of abbreviations

1 Supervisory Control And Data Acquisition
2 Common Information Model
3 Common Grid Model
4 Application Programming Interface
5 Model Context Protocol
6 Large Language Model
7 Retrieval-Augmented Generation

1 More info in our press release: https://www.iosb-ast.fraunhofer.de/de/presse/pressemitteilungen/gridcompanion_ki_agentenloesung_netzbetrieb.html.

Fraunhofer-Institut für Optronik, Systemtechnik und Bildauswertung IOSB
info@iosb.fraunhofer.de
www.iosb.fraunhofer.de

Fraunhofer IOSB Karlsruhe
Fraunhoferstraße 1
76131 Karlsruhe
Telefon +49 721 6091-0

Fraunhofer IOSB Ettlingen
Gutleuthausstraße 1
76275 Ettlingen
Telefon +49 7243 992-0

Institutsteil für angewandte Systemtechnik IOSB-AST
Am Vogelherd 90
98693 Ilmenau
Telefon +49 3677 461-0
info@iosb-ast.fraunhofer.de
www.iosb-ast.fraunhofer.de

Institutsteil für industrielle Automation IOSB-INA
Campusallee 1
32657 Lemgo
Telefon +49 5261 94290-22
juergen.jasperneite@iosb-ina.fraunhofer.de
www.iosb-ina.fraunhofer.de

Zweigstellen & Kontaktbüro

Forschungsgruppe IT-Sicherheit für Kritische Infrastrukturen ENERGIE UND WASSER
Wilhelmsplatz 11
02826 Görlitz

Fraunhofer-Zentrum für die Sicherheit Sozio-Technischer Systeme SIRIOS
c/o Fraunhofer FOKUS
Kaiserin-Augusta-Allee 31
10589 Berlin

Fraunhofer-Forschungsgruppe Smart Ocean Technologies SOT
Alter Hafen Süd 6
18069 Rostock

Forschungsgruppe Aktive Laserfasern
Moritz-Hensoldt-Straße 10
73447 Oberkochen

Kontaktbüro Beijing
Unit 0602G, Tower D1
DRC Liangmaqiao Office Building
19 Dongfangdonglu, Chaoyang District
100600 Beijing, PR China
muh@fraunhofer.com.cn

ePaper issue