AI Transforming Manufacturing:

Examples of AI & Multi-Agent Systems Powering the New Era of Generative 3D Design

Executive Summary

Industrial manufacturers face a growing gap between customer expectations for rapid, custom designs and the realities of slow, expert-dependent design workflows. Traditional order processing is constrained by manual RFQ interpretation, CAD modeling bottlenecks, and fragmented data handoffs. Recent advances in **Generative AI for 3D Design** – notably Text-to-CAD and LLM-based scripting – have automated parts of this workflow. Yet, true transformation comes with **multi-AI agent architectures**, where AI agents specialize in tasks such as simulation, manufacturability analysis, and PLM/ERP integration.

This paper outlines the technical structure of these systems, from iterative refinement loops to MCP-based interoperability, and showcases practical applications in aerospace, additive manufacturing, and composite tooling. Real-world case studies from **Siemens, Rolls-Royce, Intellegens**, and **Plyable** demonstrate dramatic gains in speed, cost, and design quality. For manufacturing leaders, the message is clear: The Multi-Agent Factory marks the next evolution in industrial intelligence, and the companies that embrace it will set the standard for manufacturing in the AI decade.

1. The Custom Manufacturing Challenge: Order Processing Bottleneck

For modern industrial manufacturers, especially those dealing with high-mix, low-volume custom orders, the traditional design process has become a major drag on profitability and speed. While customer expectations demand agility, internal operations often remain fragmented and slow, trapped in a cycle of manual repetition and data friction.

The Pain Points Eating Into Margins

The core challenge lies in converting a customer's Request for Quote (RFQ) or textual requirements into a validated, manufacturable 3D model. This process is burdened by three critical bottlenecks:

- **Extended Lead Times and Cost:** New, complex custom designs routinely take **months** to finalize. This delay is a direct result of handoffs between several domains specialists for the steps in the process, such as information gathering, requirements definition, manual CAD modeling, and simulation runs. These delays dramatically slow time-to-market and inflate costs, undermining return on investment.
- **Overloaded Expert Dependence:** The burden of custom design falls onto a small number of highly skilled CAD designers and domain experts. This creates an immediate **bottlenecked workflow** where expert availability directly limits throughput and scalability. When these individuals are tied up with information gathering or routine modeling and data manipulation, they cannot focus on their true value add.
- **Fragmented Data Handoffs:** Once a 3D model is created, the design data becomes siloed in a way. Non-engineering stakeholders, such as sales and procurement often cannot easily view, measure, or annotate the complex CAD files without purchasing costly, specialized licenses and having the necessary training to use CAD tools. This friction requires engineers to constantly export screenshots or generate static PDFs, adding additional workload with little value add.

2. Generative AI for 3D Design: Where We Are Now

Generative AI first entered the industrial design space with single-agent models focused on automating specific, initial steps of the geometry creation process. These early systems demonstrated the core premise: converting natural language or simple constraints into (digitally) tangible 3D structures.

The Rise of Text-to-CAD

Today, a single model agent, powered by a Large Language Model (LLM) and generative machine learning (ML) algorithm can be instructed to produce complex geometry. This is the era of **Text-to-CAD**, allowing engineers to describe a desired component in plain English and receive an editable, 3D model (B-Rep geometry) in return.

- **Zoo Text-to-CAD:** Companies like Zoo Design Studio offer experimental Text-to-CAD systems that generate precise B-Rep geometry directly from a text prompt (ZOO Technologies, 2025).

- **LLMs Scripting CAD:** Research has successfully demonstrated connecting LLMs, such as GPT-4, to open-source CAD platforms like FreeCAD. Given a design description, the LLM writes a CAD script, which is then executed, reviewed, and refined iteratively by the model itself (Zhang, Li, & Wang, 2024).

The Architecture and Workflow of LLMs Scripting CAD

Although there are many ways to achieve the same goal, the technical implementation of GPT-4 and FreeCAD connection reveals sophisticated multi-component systems:

Iterative Refinement Loop: The framework leverages the GPT API for script generation and LangChain for prompt management and iterative refinement, enabling the system to interpret design descriptions, generate scripts, and refine them based on execution feedback in a loop until a valid 3D model is generated. The workflow begins with natural language input which is processed by the LLM to generate a FreeCAD script, then the script is executed in FreeCAD's headless mode, and errors are fed back into the refinement loop.

Probabilistic Code Generation: LLMs such as GPT-4 operate as probabilistic generators of text, with the objective being to maximize the probability of generating syntactically correct and semantically meaningful Python scripts that accurately reflect the user's design intent. The system uses structured prompts that provide constraints, best practices, and FreeCAD-specific scripting requirements.

These achievements are profound, yet single-agent models have a crucial limitation: they excel at **geometry generation** but lack comprehensive **workflow understanding.** They are a powerful tool but not a full team; they do not natively handle the complex, sequential tasks of simulation, manufacturability checking, material costing, or integration with the PLM/ERP environment.

3. The Multi-Agent Revolution: Mimicking the Engineering Team

The next leap forward is the transition from single, isolated tools to **Multi-Agent AI Systems,** which have been made possible by recent advancements in Model Context Protocol (MCP) for CAD. This approach mirrors the real-world engineering team, where specialized engineers collaborate to create a design.

In a multi-agent system for 3D design, multiple AI agents, each trained and optimized for a specific domain, work together, passing data and decisions back and forth until the goal is achieved.

The Anatomy of a Design Agent Workflow

Below we outline a possible multi-agent system based on Xu et al. (2025)‘s model for industrial design. Instead of one “Jack-of-all-trades” model, the system divides labor among focused, purpose-built agents:

1. **RFQ Analysis Agent:** Interprets vague or fragmented customer requirements, standardizes input data, and initiates the design chain. This agent may even communicate back to the client for necessary clarifications.
2. **Styling/Aesthetic Agent:** Generates high-resolution renderings and aesthetic variations based on customer or brand requirements.
3. **CAD Agent (B-Rep Synthesizer):** Converts the conceptual design and specifications into precise, editable 3D B-Rep geometry, referencing internal design libraries and constraints.
4. **Simulation & Analysis Agent:** Automatically generates CFD (Computational Fluid Dynamics) meshes or FEA (Finite Element Analysis) models and runs complex physics simulations to predict performance (e.g., aerodynamic efficiency, structural load).
5. **DFM/Costing Agent (Design for Manufacturing):** Checks the resulting model against factory constraints, material properties, and vendor prices, providing real-time cost estimates and manufacturability feedback.
6. **PLM/ERP Integration Agent:** Ensures the final, validated design is correctly uploaded into the Product Lifecycle Management (PLM) and Enterprise Resource Planning (ERP) systems, triggering material procurement and production scheduling.

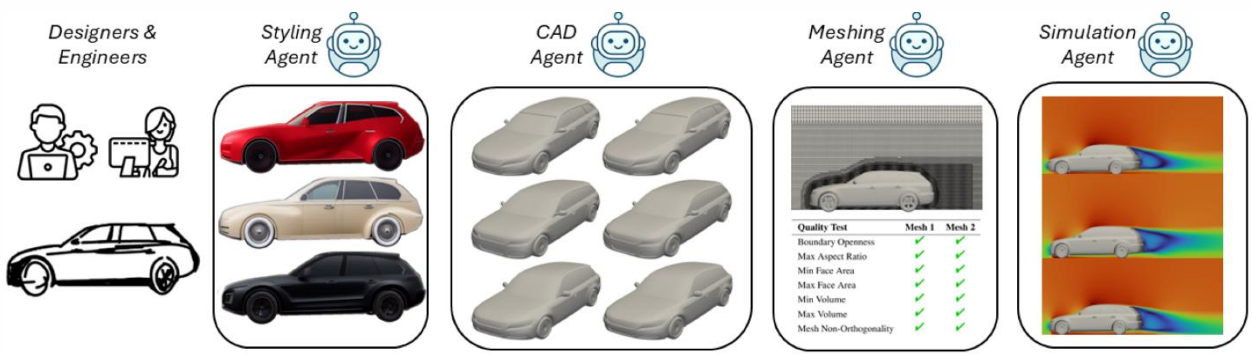

*figure 1: The multi-agent system for effective car design process. (ResearchGate.com). *preprint*

Model Context Protocol (MCP): Technical Deep Dive

What is MCP? MCP is an open protocol that standardizes how AI applications connect to various tools and services. The Model Context Protocol provides a two-way connection for AI-LLM applications to interact directly with external data sources, using a structured JSON-RPC format.

MCP servers provide functionality through three primary building blocks:

- **Tools** (software functions that an LLM can directly call to gain access to external resource data),
- **Resources** (passive data sources that (usually) provide read-only access to data for context such as file contents or database schemas), and
- **Prompts** (pre-built instruction templates that tell the model how to work with specific tools and resources)

4. Real-World Multi-Agent Implementations

While the core functionality is provided by a number of providers like AWS, n8n, IBM, etc., major manufacturers are applying principles of multi-agentic AI to achieve monumental gains in product development.

Case study: Aerospace Component Innovation: Siemens and Rolls-Royce

A prominent demonstration of GenAI's utility in 3D design was presented by Siemens and Rolls-Royce at EMO 2025, showcasing an AI-driven design-to-manufacturing process for aerospace components, specifically optimizing a lubrication and scavenge pump (Siemens Digital Industries Software, 2025).

The core application utilized an **AI-powered digital thread**, integrating topology optimization and advanced simulation within a comprehensive digital twin environment. This unified software ecosystem allowed Rolls-Royce to rapidly explore and evaluate countless design and process variations. The system ensured that designs maintained compliance with standards throughout the iteration process.

figure 2: EMO 2025 showcase of the Rolls-Royce oil pump for aerospace optimized with GenAI(siemens.com).

The results were substantial and verifiable: the AI-generated and optimized component was **25% lighter, 200% stiffer, and met a safety factor of 9** relative to the original concept. The value of this generative model extended far beyond the initial geometry. Because the generated design was inherently data-rich, it enabled downstream AI-powered tools, such as the NX CAM Copilot, to slash computer-aided manufacturing (CAM) programming time by up to 80%. This implementation demonstrates that ROI of industrial GenAI is not limited to initial creation but lies in the **data maturity of the organization** that enables automation across the entire value chain, linking design, simulation, manufacturing, and QA.

Architecture

The core implementation is set up on the **Siemens Xcelerator** business platform. The architectural backbone requires an end-to-end digitalization workflow that heavily utilizes **digital twin technology** and **simulation-driven design**. Within this workflow, three core tools define the implementation stack: **Teamcenter X** serves as the Product Lifecycle Management (PLM) and data backbone, ensuring seamless collaboration and data integrity throughout the lifecycle by utilizing characteristics and Product Manufacturing Information (PMI) data. **NX** operates as the unified CAD/CAM environment where the advanced design, simulation, and manufacturing processes are executed

The solution begins with foundational techniques such as **generative design** and **topology optimization**, which are necessary to produce breakthrough part designs for enhanced performance. The critical multi-agent integration occurs directly within the design environment through the AI-enabled copilot integrated into NX Designcenter, powered by the Microsoft Phi-3 Industrial Copilot. The copilot functions as a sophisticated Natural Language Understanding (NLU) agent, enabling users to interact with complex CAD data conversationally. Users can prompt the copilot to analyze existing design data, generating personalized and context-aware recommendations. Furthermore, the agent can execute commands to produce new design data or make critical design changes

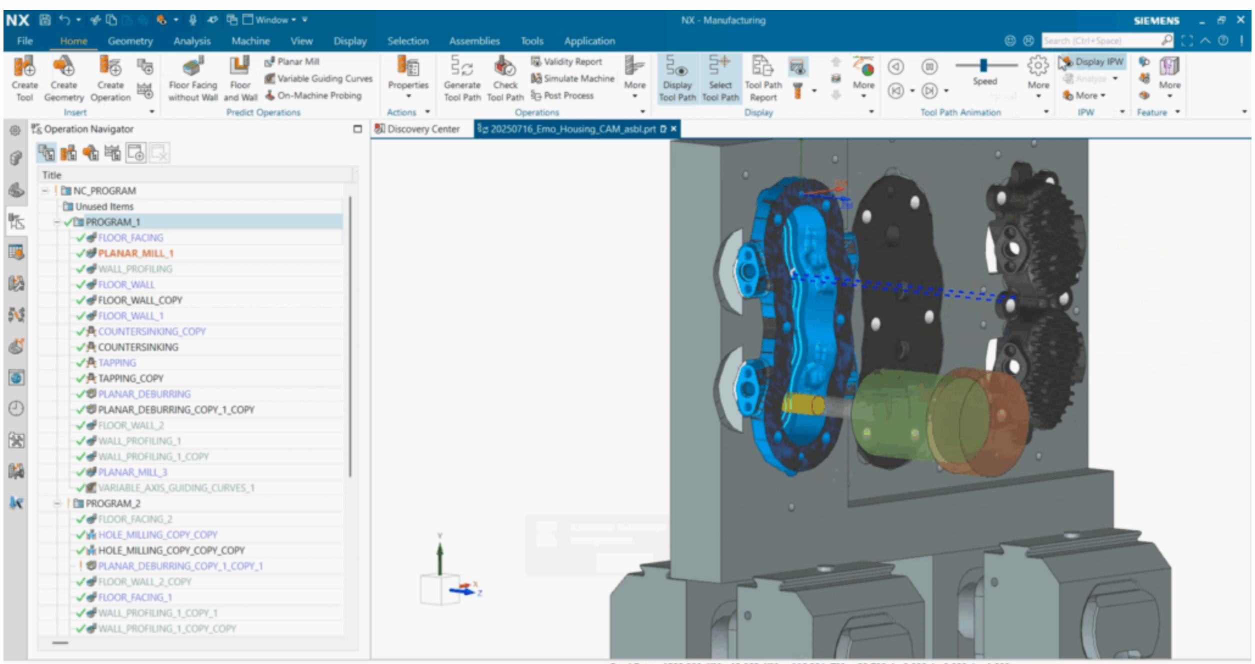

figure 3: NX CAM Copilot UI (siemens.com).

Following the design phase, a highly specialized agent, the NX CAM Copilot, takes over the manufacturing preparation workflow. This agent focuses exclusively on streamlining the time-consuming process of defining machining strategies.

The underlying methodology involves the copilot analyzing the geometric features of the component and, with minimal interaction (often just a couple of clicks), automatically generating **multiple, intelligent machining strategies**. While initially focusing on 2.5-axis and 3-axis machining, this approach aims to provide manufacturing engineers with customized, intelligent options that they can refine against company standards.

Case study: GenAI for Dynamic Quality Assurance: Real-Time Material and Process Optimization

In a collaborative project reported in a May 2025 white paper, Intellegens, Photocentric, and AMFG applied the Alchemite™ machine learning algorithm to automate and optimize advanced 3D printing processes (Intellegens, 2025).

figure 4: Production stage combines technology from Photocentric, Intellegens, and AMFG (AI-powered Additive Manufacturing).

This application of GenAI is crucial because the Alchemite™ algorithm is optimized to handle the typically sparse and noisy real-world experimental and process data inherent in materials R&D. By coupling this powerful algorithm with an 3D printer, the system is able to adjust its printing settings in real time based on process data, maximizing resource efficiency and ensuring consistency.

The successful deployment of this AI-driven approach resulted in a demonstrable reduction in waste, decreased energy usage, and lower costs. This capability directly addresses the historical challenges of inconsistency and high material cost associated with 3D printing, which is necessary to achieve the transition to high-volume, reliable serial production predicted for 2025. This signifies that AI's role in 3D manufacturing has expanded from conceptualization to dynamic process quality assurance.

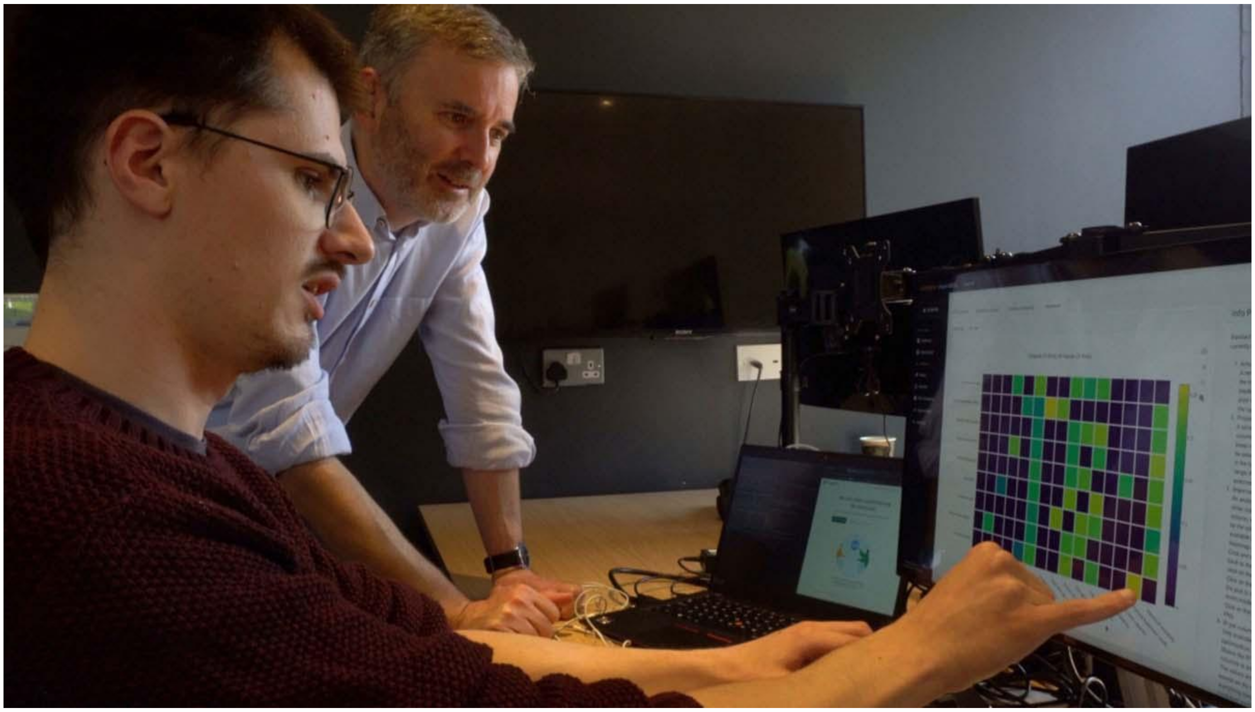

figure 5: Intelligens model training (businessweekly.co.uk)

Case Study: Tooling Design Revolution: Automating Composite Mold Generation

The composite tooling industry, critical for aerospace, marine, and high-performance automotive parts, has historically been a poster child for challenges connected to small-series custom production (lengthy lead times / high costs due to manual processes). In late 2024, Plyable introduced an AI-driven platform specifically designed to challenge these bottlenecks (Plyable, 2024).

The platform employs AI to automate the historically slow, costly, and labor-intensive mold design process, streamlining complex manual procedures involved in creating composite tooling. This approach is designed to drastically reduce costs and shorten lead times for the creation of high-precision composite molds, enhancing efficiency and competitiveness within the specialized manufacturing supply chain.

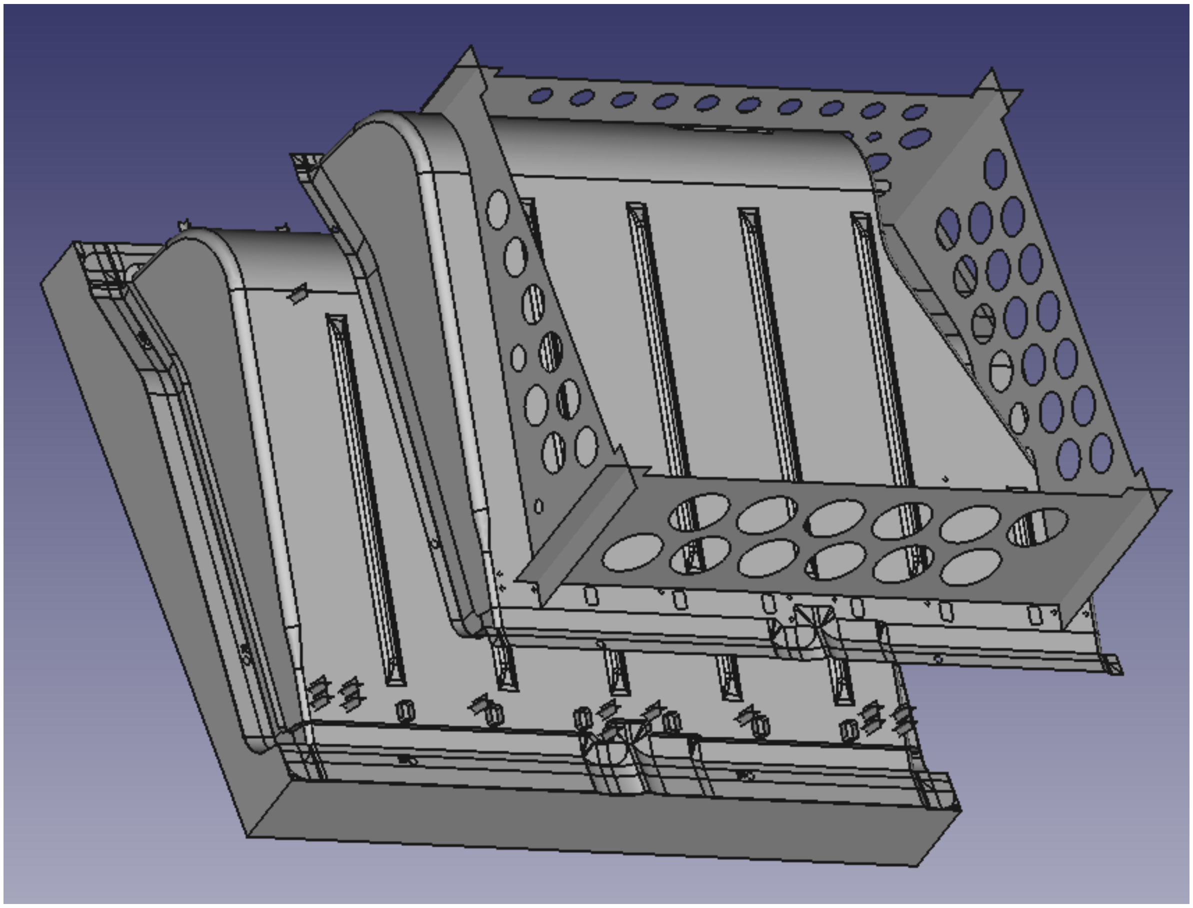

figure 6: AI and Digitisation in Composite Manufacturing (plyable.com)

The technical foundation of the platform is a combination AI and machine learning techniques. The company employs convolutional neural networks (CNNs), which are deep learning algorithms that analyze input images to assign importance to various aspects and objects, then differentiate between them. For example, the Automated Molding technology enables users to upload component design STEP files into the computational geometry agent system, which then uses agents to compute pull directions, automatically extrude shape boundaries, and add elements like laser targeting holes to generate precise, supplier-ready STEP files in minutes through a fully computer-driven process. It is not far-fetched to enrich the models with additional agents and tools to provide broader, more comprehensive implications for the manufacturer. For example, a pricing agent can calculate tooling costs based on component geometry and library of priced parts or prices from public sources, continuously improving over time and adapting to market conditions and seasonality.

These deep learning models can provide insights that trained engineers typically recognize intuitively but which are difficult to automate using traditional computational methods, particularly for tasks like geometry similarity analysis and feature identification where CNNs show excellent results

BEKO Solutions

5. Challenges & Governance: Ensuring Responsible Automation

While the potential for speed and cost reduction is massive, deploying multi-agent systems requires proactive planning to address key challenges in validation, integration, and governance.

Validation and Creativity

The core question remains: how do we trust the output?

- **Human-in-the-Loop:** AI can generate, but **human engineers must still approve and validate** the designs. The engineer's role shifts from a drafter to a strategic overseer, verifying that the AI's solution meets all technical and safety requirements.
- **The Creativity Challenge:** Can AI handle truly novel requirements and requests, or is it limited to patterns it knows well? Currently, most useful designs are simple and optimize existing designs, so manufacturers must build internal frameworks to test the **creative limits** of their deployed agents.

Integration and Data Access

Multi-agent systems cannot operate in a vacuum. They must communicate with the existing digital infrastructure and organizational setup. This is often the most complex aspect of implementation.

- **Integration:** Success hinges on the AI agents' ability to seamlessly call, access, and write back data to your existing **CAD, PLM, and ERP systems.** Not only does this require API connections, but the entire setup should be designed with integrations in mind to avoid building a parallel universe which is more burdensome than valuable.
- **Democratizing Design Review:** Using generative AI may reduce reliance on engineers for simple jobs, but even the best GenAI output is useless if non-engineers can't review it. Visualization tools, such as **SpinFire**, are crucial here. They allow sales teams, project managers, and procurement officers to view, measure, and annotate complex 3D models without requiring expensive, high-end CAD licenses, eliminating the final handoff bottleneck.

Governance, IP, and Compliance

This new technology demands a proactive approach to risk and compliance.

- **Cybersercurity & Intellectual Property (IP) Risks:** Manufacturers must establish clear policies for cybersecurity and data governance to ensure crucial IP does not leave their control, and that maximum cyber robustness is ensured. This often points toward the need for **sovereign AI**—models running on a private, secure infrastructure to protect core design IP.
- **Regulatory Compliance (EU AI Act):** As AI systems take on design responsibilities, they fall under increasing regulatory scrutiny. Leaders must ensure AI-generated designs and the underlying system adhere to global standards, including the **EU AI Act**, such as transparency and traceability of models.

6. Roadmap for Implementation: Practical Takeaways

The shift to the Multi-Agent Factory requires a strategic, phased approach. Here is a practical checklist for manufacturing leaders looking to capitalize on this transformation:

1. **Diagnosis** - Map Workflows and Bottlenecks: Conduct a forensic audit of your current custom-order process. Identify the exact steps (e.g., RFQ interpretation, information gathering, meshing, initial CAD draft) that consume the most time and are reliant on scarce experts.
2. **Experimentation** - Start with Pilot use cases: Choose a contained, high-impact process with medium complexity (e.g., automating a specific component family, or structural analysis for a subset of materials). Aim for a quick win to build learning, organizational momentum and expertise.
3. **Trust & Reliability** - Build Validation Frameworks: Define clear, measurable test protocols for AI-generated outputs. Create mandatory human review and approval gates for important steps, such as manufacturability, cost, and compliance before any design moves to production.
4. **Integration** - Plan for Integration, Not Replacement: Your goal is to augment your existing ecosystem. Ensure new AI agents are designed to communicate bi-directionally with your established CAD, PLM, and ERP platforms. Avoid siloed AI tools.
5. **Talent Augmentation** - Empower Human and AI Collaboration: Train your engineering team to work with new AI tools. Make sure they are involved from the very beginning to understand how the users will be defining constraints, managing the multi-agent workflow, and validating results.

Conclusion

The Multi-Agent Factory is no longer a futuristic concept; it is the inevitable outcome of digital transformation in manufacturing. By orchestrating specialized AI agents, manufacturers can match customers' accelerating demand for customization and enable rapid, autonomous manufacturing design without sacrificing speed or resources. This new paradigm empowers engineers to do more with less.

The time to move beyond simple process automation and into strategic multi-agent AI deployment is now, as evidenced by real-world success stories from industry leaders. By adopting a holistic, step-by-step approach to integrating AI, business leaders can unlock unprecedented speed, reduce operating costs, and secure a competitive advantage in the decade ahead.

We invite you to engage with us to understand how multi-agent workflows can be custom-tailored to your specific bottlenecks, existing infrastructure, and core compliance needs. Let's discuss accelerating your order-to-production cycle today.

About the Author

Jan Vincent Mertelj serves as Country Manager of BEKO Solutions Slovenia, where he leads the strategic expansion of innovative digital transformation solutions across the Slovenian market. With experience in telecommunications, fintech, and industrial sectors, Jan specializes in delivering software, data & AI solutions with a consultative approach. At BEKO Solutions Slovenia, he oversees the development and delivery of custom AI solutions, including RAG systems, multi-agentic AI implementations, and AI process optimization.

Jan holds a Master of Science in Strategy, Innovation and Management Control from Vienna University of Economics. His expertise combines technical innovation with strategic business acumen to deliver practical AI solutions for enterprise environments.

References

Accio. (2024). *Factory machine trends: The changing landscape of smart manufacturing.* Retrieved from https://www.accio.com/business/factory-machine-trends

aPriori. (2024). *Eaton's generative AI cuts product design time by 87%* [Case study]. Retrieved from https://www.apriori.com/resources/case-study/eatons-generative-ai-cuts-product-design-time-by-87/

Autodesk. (2024). *Generative design and generative AI: The future of manufacturing innovation.* Retrieved from https://www.autodesk.com/design-make/articles/generative-design-and-generative-ai

Intellegens. (2025). *AI-powered additive manufacturing.* Retrieved from https://intellegens.com/ai-powered-additive-manufacturing/

Plyable. (2024, October). *AI and digitisation in composite manufacturing.* Retrieved from https://www.plyable.com/2024/10/ai-and-digitisation-in-composite-manufacturing/

Precedence Research. (2024). *Generative AI in manufacturing market – growth, trends, and forecasts 2024–2032.* Retrieved from https://www.precedenceresearch.com/generative-ai-in-manufacturing-market

Siemens Digital Industries Software. (2025, September). *Taking flight at EMO 2025: Advanced aerospace design-to-manufacturing with AI and cloud innovation.* NX Manufacturing Blog. Retrieved from https://blogs.sw.siemens.com/nx-manufacturing/taking-flight-at-emo-2025-advanced-aerospace-design-to-manufacturing-with-ai-and-cloud-innovation/

YouTube. (2025). *The multi-agent factory: AI's role in the future of 3D manufacturing design* [Video]. Retrieved from https://www.youtube.com/watch?v=-djksr95QGs

ZOO Technologies. (2025). *Text-to-CAD: AI-generated 3D models from natural language prompts.* Retrieved from https://zoo.dev/text-to-cad

Zhang, Y., Li, C., & Wang, F. (2024). *Generative AI for CAD automation: Leveraging large language models for 3D modelling. Journal of Computational Design and Engineering*, 11(2), 134–149. https://doi.org/10.48550/arXiv.2508.00843

Xu, T., Müller, L., & Chen, H. (2025). *AI agents in engineering design: A multi-agent framework for aesthetic and aerodynamic car design. International Journal of Engineering Design Systems*, 8(1), 55–71. https://doi.org/10.48550/arXiv.2503.23315