WHITE PAPER

AGENTIC AI FOR PLM-CENTRIC AUTOMATION: A BLUEPRINT FOR RAPID SOLUTION DESIGN

Abstract

Product Lifecycle Management (PLM) is evolving from a static system of record into an intelligent system of understanding that continuously interprets product knowledge across tools, processes, and lifecycle stages. Rather than simply managing data or workflows, modern PLM acts as a product intelligence layer, integrating CPQ, PLM, ERP, MES, and field systems into a unified enterprise memory.

Manufacturers face growing product complexity, variant proliferation, and stringent compliance requirements. Fragmented systems and manual reconciliation often result in configuration drift, rework, and traceability gaps.

This Paper Introduces a PLM-centric Agentic AI architecture composed of specialized agents governed by defined policies. These agents automate key processes, including requirements capture, eBOM → mBOM → BOP transformation, engineering change management (PR/ECR/ECN/MCO), work instruction generation, and ERP/MES synchronization, while preserving serialized field configuration integrity.

Each agent leverages retrieval-augmented reasoning (RAG) over engineering, service, and compliance data to produce grounded, explainable outputs. Operating continuously, agents subscribe to PLM events and return structured insights such as BOM deltas, routing recommendations, and work instruction drafts.

The outcome is elimination of configuration drift, end-to-end digital thread traceability, and predictable engineering-to-operations cycles, significantly improving productivity and enabling intelligent lifecycle orchestration.

Infosys
Navigate your next

Introduction

Manufacturers across sectors face rising complexity: expanding product portfolios, dynamic customer configurations, stricter regulations, and volatile supply chains. These pressures demand faster, more precise, and more resilient engineering and manufacturing operations. Traditional PLM systems—designed primarily as systems of record—are no longer adequate.

Modern operations require PLM to evolve into a system of understanding: one that interprets engineering intent, reasons across multidomain constraints, and maintains continuous alignment with CPQ, PLM, MES, ERP, and field systems. Yet, Most organizations still depend on manual interpretation, disconnected workflows, and fragile integrations. This results in delayed visibility into feasibility and cost, inconsistent eBOM→mBOM→BOP transformations, slow ECN propagation, outdated work instructions, and configuration drift across design, factory, and field—leading to cycle time variation, rework, and compliance risks.

Agentic AI provides a Fundamental shift. Instead of embedding logic in static workflows, specialized AI agents handle requirements interpretation, BOM/BOP derivation, change orchestration, digital WI generation, and cross-system synchronization. Operating over PLM as the authoritative backbone, these agents collaborate, reason over structured and unstructured data, and maintain a continuously coherent digital thread.

This white paper introduces a PLMcentric Agentic AI architecture designed to achieve zero configuration drift, full lifecycle traceability, and predictable delivery. It explains how multiagent orchestration eliminates interpretation gaps, accelerates engineering throughput, and enables realtime synchronization across CPQ, PLM, MES, ERP, and field systems—supported by early deployment outcomes.

Present Challenges and Problem Statement

Despite investments in PLM, MES, ERP, and CPQ, most manufacturers still face misalignment between design intent and actual production or field outcomes. The gap arises from rising product complexity, supplychain volatility, and the tight interdependence across engineering, manufacturing, and service data. Traditional PLM workflows—manual, static, and interpretationheavy—can no longer keep pace.

1.Fragmented lIfecycle Data Across CPQ, PIM, MES, ERP, and Field

- Core teams rely on systems that cannot interpret each other's data.
- CPQ options rarely map cleanly to engineering-valid configurations.
- eBOMs and CAD live in PLM: routings in MES: cost/availability in ERP.
- Field configuration data is delayed or incomplete.

Impact: Manual reconciliation across systems increases rework, delays, and cycle time unpredictability

2. Manual eBOM → mBOM → BOP Transformations

Organizations depend on a few experts to convert engineering structures into plant-specific mBOMs and process plans.

- Rules and tacit knowledge live in people, not systems.
- Capability and sequencing knowledge is dispersed.
- Plant effectivity is applied manually.

Impact: Structural bottlenecks, variable cycle times, and quality risks. Variant growth makes this approach unsustainable.

3. Slow, Error-Prone Change Propagation

- Engineering changes seldom flow cleanly across PLM, MES, and ERP.
- Downstream systems receive updates late or not at all.
- Routings, cost data, and procurement triggers misalign with effectivity.

Impact: Outdated instructions and wrong revisions cause scrap, rework, and compliance exposure.

4. Configuration Drift Across Design, Manufacturing, and Field

- Even with controls, systems diverge over time.
- Field units operate on revisions not reflected in PLM.
- MES runs outdated routings or tooling assumptions.
- ERP cost/item data drifts from engineering's latest state.

Impact: Traceability gaps, compliance risk, warranty cost, and downtime.

5. Digital Work Instructions Lag Behind Engineering Changes

- Instructions remain manually authored, knowledge-dependent, and updated reactively.

Impact: Operators follow outdated or incorrect guidance, creating deviations, quality escapes, and safety risks.

6. Workforce Productivity Loss Due to Manual Interpretation

- Engineers and planners spend large effort translating system outputs, reconciling spreadsheets, resolving inconsistencies, and searching for correct versions.

Impact: 25–40% of engineering capacity is consumed by non-value-added work, slowing NPI and variant releases.

Why PLM-Centric Agentic AI Architecture Is Now Core Infrastructure

PLM is no longer just a system of record — it's becoming the digital backbone for orchestrating the entire product lifecycle. With Agentic AI, PLM transforms from a passive repository into an active decision-making and execution engine that connects engineering, manufacturing, supply chain, quality, and after-sales in real time.

Below is the real reason it has become core infrastructure, not "one more technology".

1. Engineering complexity has exploded

Modern products have:

- High variant proliferation
- Multi-disciplinary engineering (mechanical, electrical, software)
- Rapid ECO/ECR cycles
- Compressed launch timelines

Traditional PLM workflows cannot cope with manual checks and slow cross-system alignment.

Agentic AI automates alignment across PLM → CAD → MES → ERP → Service, making lifecycle management scalable.

2. PLM is the only system with the full product context

AI agents need:

- Complete BOM/BOP structure
- Design rules
- Requirements
- Configurations
- Simulation data
- Change history

Only PLM has this end-to-end truth.

Without a PLM anchor, AI outputs become unreliable or non-traceable.

Thus, **PLM becomes the "source-of-grounding" for all AI reasoning.**

3. Agentic AI turns PLM into an autonomous operations coordinator

Agentic AI can:

- Auto-compose eBOM → mBOM → BOP
- Detect inconsistencies between design & manufacturing
- Generate work instructions
- Trigger purchase requests or ECOs
- Coordinate MES & supply chain alignment
- Simulate manufacturing impacts instantly

This shifts PLM from records to **real-time orchestration**, making it core infrastructure for digital factories.

4. Enterprises need closed-loop traceability for compliance, warranty, and sustainability

Today's regulatory pressure demands:

- Full lineage from design → manufacturing → service
- Auditable change traceability
- Digital thread continuity

Agentic AI can maintain this digital thread automatically, closing the loop across:

- Engineering
- Quality
- Service

This gives organizations **risk-free change management** — a non-negotiable requirement globally.

5. Cloud-native PLM platforms made AI integration frictionless

Modern PLM platforms (e.g. 3DX, Windchill+, Teamcenter X, Aras) are:

- API-first
- Graph-based
- Model-driven
- Cloud-deployable

This makes them ideal foundation for:

- LLM-based reasoning
- RAG over engineering data
- Autonomous agent orchestration
- Digital thread intelligence

Thus, PLM becomes the **nerve center for enterprise AI automation.**

Enterprise Solution:

Traditional rule-based automation cannot handle modern engineering complexity, unstructured data, dynamic conditions, or cross-system misalignments. Agentic AI can reason, adapt, retrieve context, validate with rules, detect drift, and orchestrate PLM-ERP-MES workflows automatically — enabling zero drift, real-time synchronization, and predictable delivery.

Achieving End-to-End Digital Continuity Across CPQ → PLM → MES → ERP → Field

Enterprises can break longstanding fragmentation across engineering, manufacturing, and service functions by deploying a PLMcentric Agentic AI architecture that establishes continuous, closedloop digital continuity across all lifecycle systems. Instead of relying on human interpretation and manual synchronization, the solution activates a coordinated network of specialized AI agents that read, reason, and act on lifecycle data in real time. This creates a self-correcting digital thread that maintains alignment from configuration and design through manufacturing execution and field operation.

The proposed solution transforms PLM from a static repository into an intelligent product backbone—a system that continuously contextualizes requirements, evaluates constraints, propagates changes, and verifies configuration integrity throughout the enterprise.

A coordinated set of AI agents capture requirements from CPQ/PLM/MES, autocompose eBOM→mBOM→BOP, orchestrate PR (purchase request) /ECR/ECN/MCO, generate digital work instructions, synchronize changes across ERP/MES, and autodeliver final configuration/parameter files to field units—ensuring zero drift, full traceability, and predictable delivery."

Primary business outcomes:

- 100% change propagation across PLM⇄MES⇄ERP⇄CPQ
- Variant-correct mBOM/BOP and work instructions on every order
- Automated SKU/parameter pack generation to field units
- Audit-ready (PR/ECR/ECN evidence) and learn-from-operations feedback loop

PLM Centric Platform & technical components

Below is the proposed architecture view of a cloudnative PLM platform powered by modular technical components—microservices (semantic search, agentsmanagement, blob storage, BOT interactive, etc.), APIs, event buses, a multiagent system, and governed product data—to drive integrated, autonomous lifecycle operations.

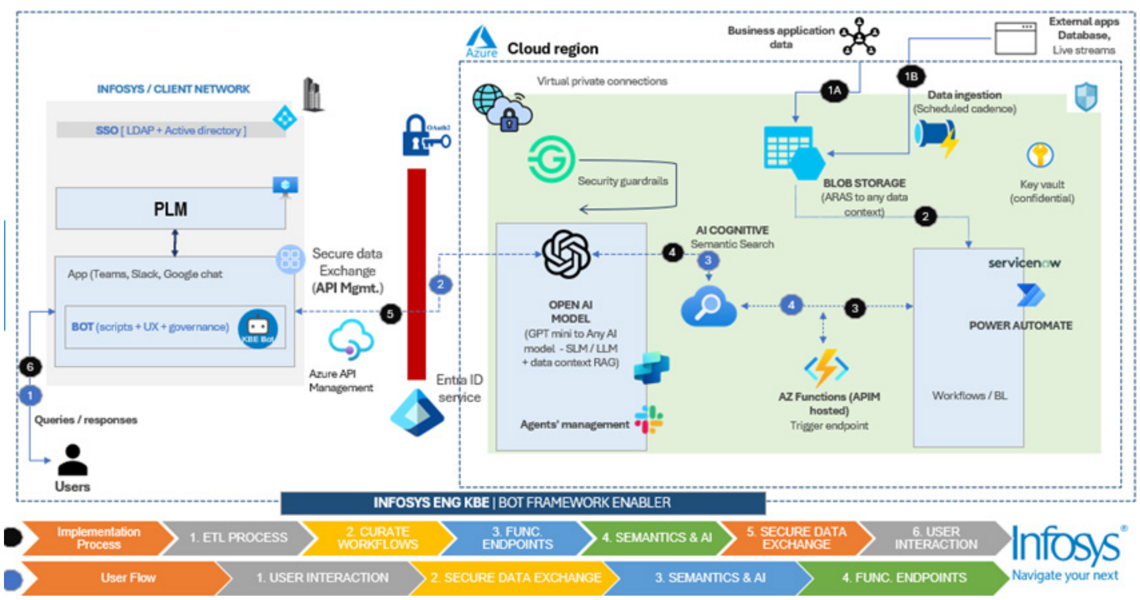

In this architecture, users submit queries through enterprise channels like Teams or Slack, which are securely routed from the client network to Azure through API Management and authenticated using Entra ID. PLM data from (PLM platform) — along with complementary business data from CPQ, MES, and ERP—is periodically ingested into Azure Blob Storage and indexed through Cognitive Search to form a unified lifecycle intelligence layer. When a request arrives, **AI agents** retrieve the right context across PLM configurations, CPQ options, MES routings, and ERP cost/availability data, then use OpenAI models to interpret intent, validate constraints, and determine what action to take. Azure Functions orchestrate these agent-driven decisions—such as generating insights, checking manufacturability, reconciling BOM/BOP discrepancies, or triggering updates across MES or ERP systems. Downstream actions, approvals, and workflows are executed via Power Automate or ServiceNow, with Key Vault and cloud security guardrails ensuring full compliance. The completed result is then returned to the user through the chatbot, delivering a secure, automated, closed-loop flow across the entire CPQ-PLM-MES-ERP digital thread.

Physical View (Deployment & Infrastructure)

Azure Cloud Region – Secure boundary	Key Vault – Secret management
Azure Functions – Compute layer	Entra ID – Identity provider
Cognitive Search – Indexing/search engine	API Management – Edge gateway
Blob Storage – Data Lake for PLM/MES/ERP/CPQ extracts	Virtual Networks + Private Endpoints – Secure data paths
	Integration with on-prem PLM/MES/ERP over VPN/VNet Peering

At a high level, the solution operates through **six collaborating AI agents**, each working in a closed-loop pattern across CPQ, PLM, MES, ERP, and Field systems to ensure zero-drift execution and full lifecycle alignment

1. Context & Subscription Agent - Handles user context, plant/program selection, and scoping. It ensures the right datasets and rules apply.

2. Planner/Orchestrator Agent - Breaks the business request into actionable tasks. Delegates work to the correct downstream agents. Core orchestration intelligence.

3. Manufacturing Definition Agent - Interprets engineering intent from PLM. Performs manufacturability checks, capability mapping, BOP validation.

4. Orchestration Agent - Executes multi-step workflows, Bridges between PLM → MES → ERP systems. Ensures correct sequence of operations.

5. Compliance & Costing Agent - Performs cost rollups, supply checks, ERP validations. Ensures compliance constraints before release (regulatory, quality, plant readiness).

6. Coordinator Agent- Gatekeeper for release. Ensures all upstream validations are completed. Manages final handover to MES/Field.

Impact Area	Executive Benefit
Engineering throughput	Faster BOM/BOP queries, ECN evaluation, design clarity
Cycle time reduction	Automated workflows eliminate manual steps
Quality & compliance	AI ensures consistency and traceability
Labor efficiency	Fewer repetitive tasks for engineers & planners
Data democratization	Anyone can query PLM data conversationally
Zero-drift operations	AI continuously aligns PLM ↔ MES ↔ downstream systems

Step-by-Step Implementation Strategy & Implementation Workflow

For Achieving PLM-Centric Agentic AI Infrastructure Integrated with CPQ, MES, and ERP

A PLM-centric Agentic AI infrastructure cannot be deployed as a single project; it requires a structured, phased implementation that builds the data foundation, agent behaviors, orchestration logic, and cross-system integrations. The strategy below outlines a practical, real-world, programmable approach followed by a factory-ready execution workflow.

Step-by-Step Implementation Strategy

Step 1:
Establish the Product Data Foundation in PLM

Objective: Create the authoritative, structured, lifecycle-controlled product core.

Activities:

- Rationalize eBOM/mBOM structures
- Clean engineering and manufacturing attributes
- Define product rules, variants, options, constraints
- Build the canonical data dictionary connecting CPQ ↔ PLM ↔ MES ↔ ERP
- Load historical product configurations and change history

Outcome: PLM becomes the single source of product truth for all downstream agent operations.

Step 2:
Identify High-Value Automation Candidates

Objective: Select processes that deliver the highest ROI when automated.

Typical candidates:

- eBOM → mBOM conversion
- BOP/Routing generation
- Digital Work Instructions
- ECR/ECO impact analysis
- CPQ configuration validation
- MES BOM / Routing synchronization
- ERP material master and production-version creation

Outcome: A prioritized automation roadmap for agent development.

Step 3:
Build the Multi-Agent Architecture

Objective: Create specialized AI agents grounded on lifecycle context.

Agents typically include:

- BOM Agent – reconciles eBOM ↔ mBOM
- BOP/Routing Agent – generates process plans
- Change Agent – propagates impacts across systems
- MES Sync Agent – aligns PLM and MES
- ERP Integration Agent – handles material masters, costing fields
- CPQ Validation Agent – checks configuration rules

Outcome: Modular agent ecosystem operating on PLM as the reasoning anchor.

Step 4:
Implement the Orchestration & Event-Driven Integration Layer

Objective: Ensure agents coordinate actions end-to-end across systems.

Key enablers:

- Event bus to publish lifecycle events (BOM release, ECO approval, CPQ order, MES deviation)
- Process orchestrator to manage agent sequencing
- Context store for agent handoff
- Error-handling and reconciliation logic

Outcome: A coordinated AI workflow where agents autonomously trigger next actions.

Step 5:
Integrate CPQ, MES, and ERP Systems

Objective: Enable closed-loop automation from customer configuration to shop-floor execution.

Integrations include:

- CPQ → PLM: variant validation, feasible configuration check
- PLM → MES: mBOM, BOP, Work Instructions, resource lists
- PLM → ERP: material cards, production versions, sourcing details
- MES → PLM: deviations, rejects, quality alerts triggering ECO proposals

Outcome: Smooth, traceable, automated data flow across the digital thread.

Step 6:
Deploy Production Use-Cases in Waves (Agile)

Wave examples:

- Wave 1: BOM/BOP automation
- Wave 2: ECO/ECR intelligent propagation
- Wave 3: CPQ configuration-to-manufacturing automation
- Wave 4: MES–PLM closed loop
- Wave 5: ERP auto-material/production-version creation

Outcome: Tangible value delivered continuously without big-bang risk.

Step 7:
Establish KPI Framework & Continuous Improvement Loop

Track KPIs such as:

- BOM accuracy
- Manufacturing readiness lead-time
- Order release cycle
- ECO cycle time
- WI update lag
- mBOM/BOP creation effort
- Rework and scrap reduction

Outcome: Data-driven evolution of agent performance and process improvements.

Business Outcomes Realized

A successful implementation yields measurable improvements:

- 40–60% reduction in order-to-manufacturing release time
- 50–70% faster BOM/BOP creation
- 70–80% shorter WI update cycle
- 30–40% reduction in ECO processing time
- 20–30% fewer production delays due to incorrect configurations
- Improved on-time delivery and first-time-right builds

To illustrate the full E2E integration flow, the scenario below walks through a complete lifecycle—from Quote generation through PLM, MES, ERP, and ultimately to Field feedback.

A) CPQ → PLM (Initiate & Author)

1. Sales engineer configures a variant (e.g. Rotor size, Hub height, Power rating, Noise mode etc.) in CPQ and gets a Config Signature (hash of selected options/constraints).
2. In chat, the user types: "Create engineering pack for Quote Q-123."
3. Bot → API Mgmt. → Function validates identity, pulls the quote, and pushes the config signature to PLM.
4. LLM (with RAG) summarizes any rule conflicts or missing parts and suggests the closest reference eBOM.

Outputs: eBOM candidate, gap list, draft ECR.

B) PLM (Enrich & Create)

1. Engineering confirms/edits the eBOM, generates the mBOM and BOP, and attaches work instructions (WI).
2. Function triggers Power Automate for approvals; if deviations arise, ServiceNow change tasks are created.
3. Approved objects are versioned; snapshots (eBOM/mBOM/BOP/WI) land in Blob and get indexed.

Outputs: Released eBOM/mBOM/BOP/WI revisions + ECN.

C) PLM → ERP (Plan & Orderability)

1. The user clicks "Release to ERP" in chat.
2. Function transforms PLM objects to ERP-ready payloads (materials, BOMs, routings) and pushes through APIM.
3. ERP confirms materials, sources, costing, and readiness; status flows back to the chat with links.

Outputs: Material masters, BOMs/routings active in ERP; orderable SKU created.

D) ERP → MES (Execute)

1. Production order drops; Function publishes as-planned ops and WIs to MES.
2. MES executes; as-built/serial genealogy and station params stream to Blob and are indexed.
3. Outputs: As-built vs as-planned record; NCs, reworks; first-time-right metrics.

E) MES/Field → PLM (Closed Loop)

1. Quality events or Field telemetry raise anomalies; Function opens ServiceNow tickets and/or proposes ECRs in PLM with evidence bundles (images, logs, deltas).
2. LLM composes a root-cause summary with citations (process step, drawing note, revision) and suggests variant rule updates (feeding CPQ) or WI changes (feeding MES).

Outputs: Faster ECO cycles, updated constraints in CPQ, corrected WIs in MES, and traceability across the thread.

Proposed two steps implementation roadmap

Step 1:
Clean the Data & Make It Usable (AI-Assisted Data Management)

Goal: Make the existing PLM data easier to find, trust, and use.

What happens in this phase:

- AI enables users search for parts, documents, and changes across systems faster.
- Simple agents monitor data issues like duplicates, missing attributes, or outdated versions.
- Chat-style AI interfaces reduce complexity — users can ask questions like "Show me the latest mBOM for Variant X."
- No major workflow automation yet — focus is on "better insight, less effort."

Step 2:
Achieve True Agentic Intelligence (Full Agentic Capability)

Goal: Let AI handle complex, multi-system engineering and manufacturing workflows end-to-end.

What happens in this phase:

- Foundation: deploy event bus, schema registry, PLM read/write APIs, and a rules engine.
- Pilot agents: implement Order Intake, Configuration Resolver, and PLM Master Data agents in monitoring mode.
- Close the loop: add ERP Orchestration and MES Execution agents with HITL escalation.
- Scale and optimize: introduce RAG knowledge layer, MLOps pipelines, supplier agents, and drift remediation automation.

Risks and mitigations

- Model hallucination: mitigate by pairing LLM outputs with deterministic rules and RAG sources; require HITL for low confidence.
- Master data inconsistency: enforce PLM as canonical source and block direct MES/ERP edits without ECOs.

POC Work Breakdown Structure (WBS)

Choose one high-value, cross-system workflow that clearly showcases PLM-centric automation, e.g.:

- CPQ-driven variant resolution → eBOM/mBOM generation
- Automated BOM/BOP transformation for a single product family
- ECO impact assessment with downstream ERP/MES alignment
- Zero-drift detection between PLM BOM/BOP and MES as-built

Pilot scope = 1 product family + 1 workflow + 3–5 KPIs.

A brief work breakdown structure to implement PLM centrical integrated solution:

Define Scope & Boundaries

- Select one product family (limited parts/variants)
- Identify one POC scenario (eBOM→MBOM, ECO impact, variant resolution, WI gen)
- Define system boundaries (CPQ, PLM, ERP, MES involvement)

Success Criteria & KPIs

- Define automation %
- Cycle-time reduction target
- Accuracy & drift detection KPIs
- HITL thresholds

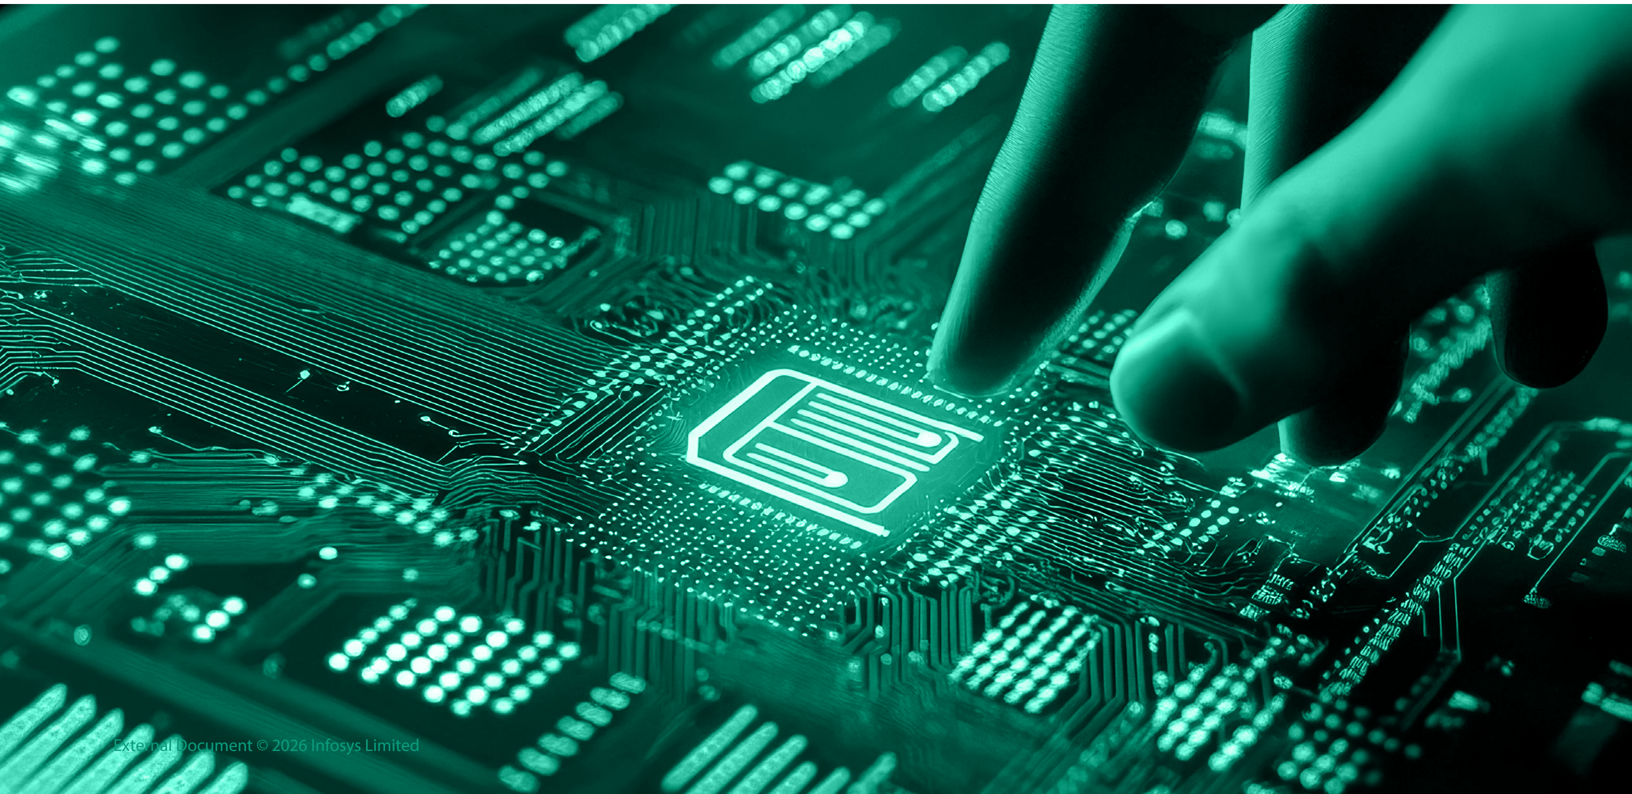

Conclusion

Deploying a coordinated set of AI agents marks a foundational shift in how modern manufacturers operate. This PLM-centric Agentic AI infrastructure eliminates traditional friction between engineering, manufacturing, and supply chain functions by grounding every action in authoritative product context. The result is a digital thread that is no longer passive or fragmented, but active, intelligent, and self-maintaining. By automating lifecycle-critical tasks with precision and traceability, organizations achieve faster order execution, zero drift between design and production, full change transparency, and dramatically improved first-time-right performance. This coordinated agent ecosystem not only accelerates delivery and reduces operational delays, but also unlocks higher productivity, predictable margins, and resilient, scalable manufacturing operations, positioning enterprises to compete and innovate with unprecedented agility.

References

1. Teamcenter Copilot (Siemens) — GenAI inside PLM: Teamcenter Copilot brings conversational search, document intelligence, BOM exploration, and visual (image-to-part) search directly into PLM, accelerating engineering productivity and grounded answers from PLM data. (https://blogs.sw.siemens.com/teamcenter/plm-genai-teamcenter/)

2. Teamcenter 2506 Copilot + GenAI (Siemens) - deployment notes Short guide/webinar on deploying, Teamcenter AI (AWS/Azure/Llama options), BOM interaction via natural language, and image-to-part capabilities—handy for POC setup. https://blogs.sw.siemens.com/teamcenter/teamcenter-2506-copilot-genai/ \[blogs.sw.siemens.com\] (https://blogs.sw.siemens.com/teamcenter/teamcenter-2506-copilot-genai/)

3. Windchill AI (PTC) - Agentic AI for the digital thread: PTC previewed Windchill AI and a Document Vault AI agent to extract & reason over PLM docs, alongside Codebeamer AI and ServiceMax AI for end-to-end digital-thread scenarios. https://www.ptc.com/en/news/2025/ptc-brings-ai-powered-plm-to-hannover-messe \[ptc.com\] (https://www.ptc.com/en/news/2025/ptc-brings-ai-powered-plm-to-hannover-messe)

4. Windchill AI product page (PTC) - parts rationalization & assistants: Describes AI parts rationalization, chat assistants over PLM data, and an agentic digital thread to automate cross-lifecycle decisions. https://www.ptc.com/en/products/windchill/windchill-ai \[ptc.com\] (https://www.ptc.com/en/products/windchill/windchill-ai)

5. Dassault 3DEXPERIENCE — AI-powered Virtual Companions (Agentic Platform): Dassault introduces Virtual Companions (Aura/Leo/Marie) on a 3DEXPERIENCE agentic platform to orchestrate requirements, engineering, and science-driven reasoning across the lifecycle. https://www.3ds.com/newsroom/press-releases/dassault-systemes-unveils-new-way-working-industry-ai-powered-virtual-companions \ (https://www.3ds.com/newsroom/press-releases/dassault-systemes-unveils-new-way-working-industry-ai-powered-virtual-companions)

6. CATIA / 3DEXPERIENCE — AI-driven generative engineering: AI in CATIA for generative experiences, topology optimization, and virtual companion interactions; ties into PLM to accelerate design-to-manufacture. https://www.3ds.com/products/catia/ai-driven-generative-experiences

7. Aras InnovatorEdge — API gateway & AI agents over the digital thread: Aras "Edge" exposes PLM via governed APIs and provides services for designing/running AI agents (discover/enrich/amplify) that act on PLM + ERP/MES data with guardrails. (https://aras.com/en/capabilities/aras-innovatoredge)

8. Aras — AI-Assisted Search & Intelligent Assistant (Azure OpenAI/Copilot): Press release on AI-assisted search across structured/unstructured PLM data and a conversational assistant integrated with Aras Innovator SaaS (Azure OpenAI). (https://aras.com/en/news/press-releases/2025/03/aras-introduces-ai-assisted-search-and-ai-powered-intelligent-assistant-for-aras-innovator-saas)

9. Cognite Data Fusion (CDF) - Industrial AI platform with PLM/MES/ERP connectors: CDF unifies OT/IT/ET (historians, MES, ERP, PLM) into an industrial knowledge graph with Generative/Agentic AI —useful for PLM-adjacent manufacturing twins and analytics. https://www.cognite.com/en/industries/discrete-manufacturing \[cognite.com\] (https://www.cognite.com/en/industries/discrete-manufacturing)

10. 10.SymphonyAI IRIS Foundry — Industrial DataOps + Agentic AI Flows: IRIS Foundry builds unified industrial knowledge graphs and agentic workflows; includes IRIS Copilot (Azure OpenAI) and connectors across IT/OT/engineering (PLM/MES/ERP). (https://www.symphonyai.com/industrial/iris-foundry/)

11. SAP Digital Manufacturing + Joule (AI copilot): Joule is integrated into SAP Digital Manufacturing, enabling conversational insights and AI functions alongside S/4HANA integration—relevant for ERP↔MES↔PLM process alignment. (https://community.sap.com/t5/technology-blog-posts-by-sap/joule-setup-for-sap-digital-manufacturing/ba-p/14211372)

12. 12. Infosys Topaz — Agentic AI Foundry (cross-platform, model-agnostic): Infosys Agentic AI Foundry provides pre-built vertical/horizontal agents, open model abstraction, governance & evaluation frameworks—ideal to host PLM-centric agents at scale. https://www.infosys.com/services/data-ai-topaz/offerings/agentic-foundry.html \[infosys.com\] (https://www.infosys.com/services/data-ai-topaz/offerings/agentic-foundry.html)

About the Authors:

Sibaji Pattanaik

Principal Consultant, IOT, Infosys

Kapil Manikrao Jadhav

Industry Principal, IOT, Infosys

Infosys
Navigate your next

For more information, contact askus@infosys.com