The context imperative: how the Smart Manufacturing Interoperability Platform can be the strategic core of agentic AI strategies

Executive briefing for manufacturing leaders

KEARNEY

100 Years of **Impact**

The structural barrier to scalable AI in manufacturing

For decades, manufacturing enterprises have sought to transform vast amounts of machine data into actionable intelligence. Yet, even in the Industry 4.0 era, the central bottleneck remains: fragmentation. Data is siloed across vendor-specific operational technology and disparate information technology stacks, each speaking a different language.

In addition, data lacks context, provenance, discoverability, and meaning. In the past, big data analytics projects were predicated on marshaling vast amounts of uncontextualized data from dozens of data sources to cloud-based data lakes and then relying on armies of data engineers and data scientists to add context and meaning. This monolithic and linear approach to AI implementation is no longer appropriate in the age of agentic AI. Putting a large language model (LLM) on top of a data lake won't work.

These structural flaws prevent the scaling of digital transformation and doom most AI initiatives to either pilot purgatory or exorbitant enterprise projects with negative ROI. These projects often involve vast amounts of bespoke engineering due to reimplementation of existing connections and the translation of incompatible data formats.

This is unsustainable. Repairing these flaws can unlock a roughly $900 billion value opportunity, representing about 5 percent (as a midpoint savings estimate) of global manufacturing value added in 2024 of $16.8 trillion (see figure 1).

Figure 1
Using AI in manufacturing can unlock a wealth of opportunities

Estimated global impact of AI on the manufacturing industry

	AI impact range	Improvement bucket	Impact[1]
Materials	— 10–20 percent reduction in quality costs — 20–50 percent reduction in inventory holding costs ~ 20 percent reduction in raw material waste	— Yield — Scrap — Rework	~1–2 percent net cost reduction = $170 billion–$340 billion annually
Equipment	— 0–40 percent lower maintenance costs — 30–50 percent downtime reduction ~ 15 percent energy cost reduction	— OEE — Asset reliability — Energy intensity	~1–3 percent net efficiency gain globally = $170 billion–$510 billion annually
People	— 7–20 percent productivity improvement (shop floor) — 45–55 percent productivity uplift in technical roles	— Labour productivity — Knowledge work efficiency — Error reduction	~1–2 percent labor cost improvement = $170 billion–$340 billion annually

[1] Based on global manufacturing value added: about $16.8 trillion (2024). Baseline manufacturing value added data points for each country (as percent of total GDP and percent split by industry) were obtained from the World Bank, 2024.

Source: Kearney analysis

This is achievable through AI-driven improvements across productivity, asset uptime, and quality while simultaneously improving process and occupational safety outcomes. Figure 2 shows an estimation of the mid-point value opportunity by industry and region.

Unlocking this opportunity requires manufacturing data with context, meaning, and discoverability. CESMII's Smart Manufacturing Interoperability Platform (SMIP) clears these structural barriers by creating a context layer: a universal semantic layer that unifies and standardizes all plant data. It also provides the essential ontology through both a knowledge graph and an object-oriented information model. This contrasts with the typical "stovepipe" architectures that are an accumulation of use case by use case applications, each justified on their own merits but are poorly integrated and developed without a comprehensive appreciation of the added value and scalability that comes from end-to-end data lineage and a common semantic foundation (see figure 3 on page 3).

AI can enhance productivity, uptime, and quality while also improving process and occupational safety outcomes.

Figure 2
Using AI in manufacturing could unlock a wealth of potential

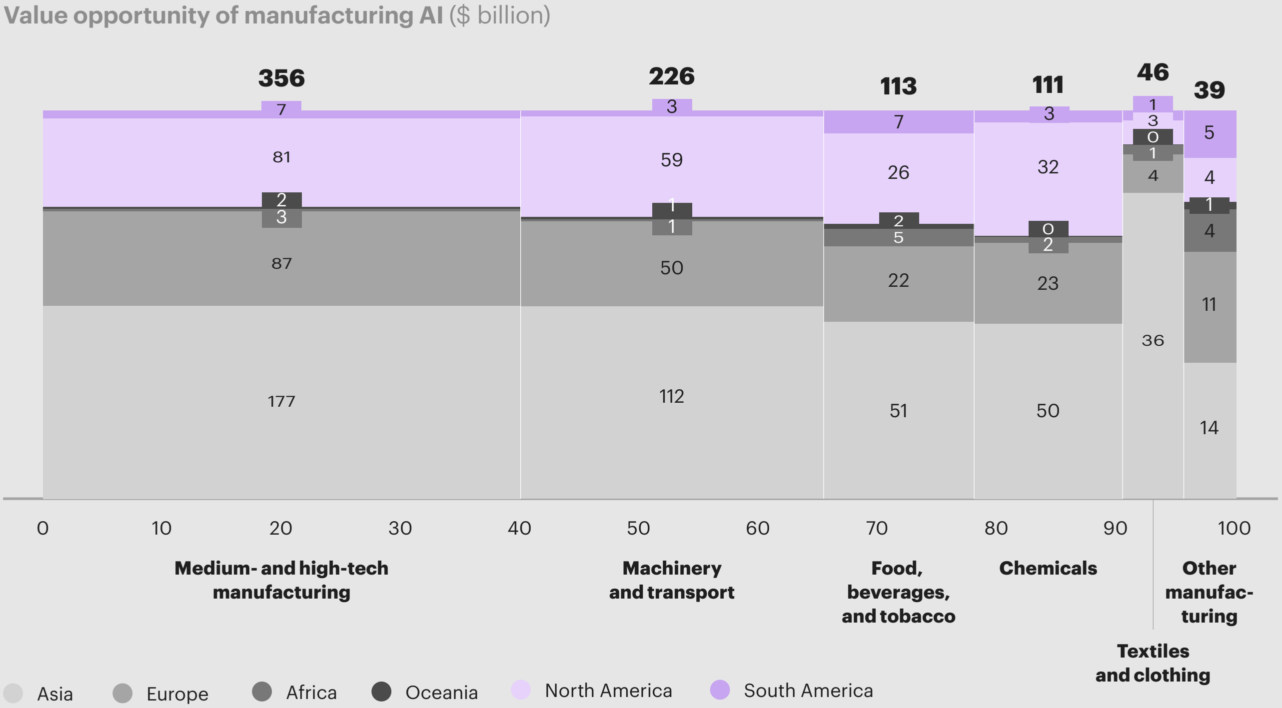

Sources: World Bank data on manufacturing value added (2024); Kearney analysis

Figure 3

Traditionally, stovepipe architectures inhibit scaling manufacturing applications

Stovepipe architectures: use case by use case and justified on its own merits

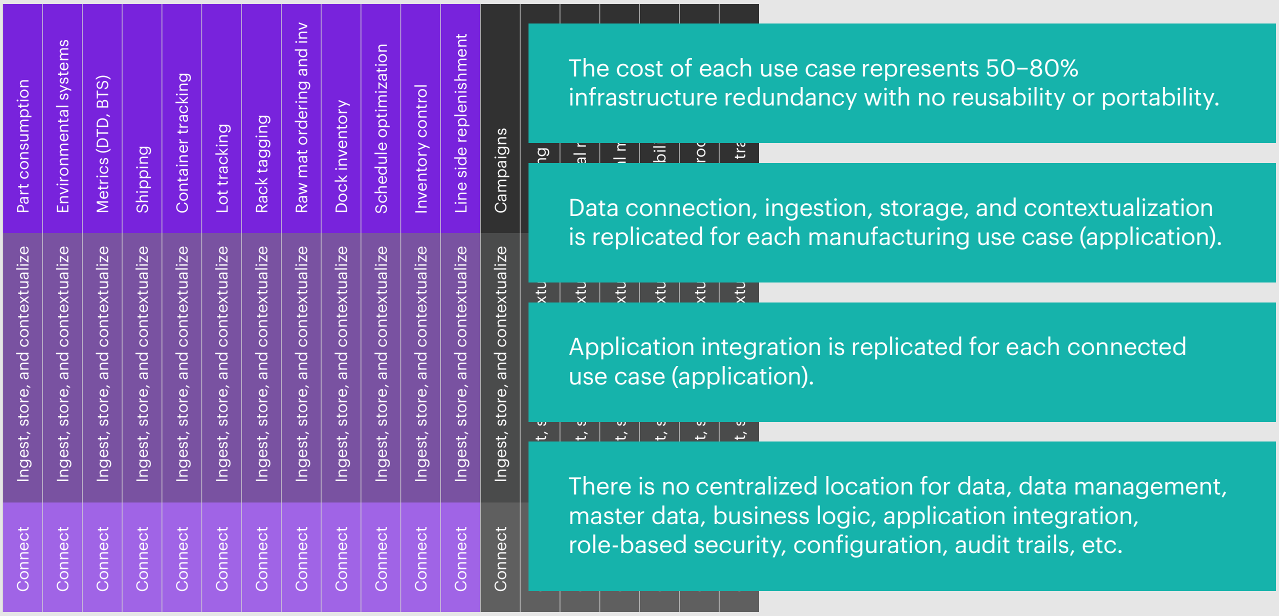

Stovepipe vs. interoperability platform cost comparison

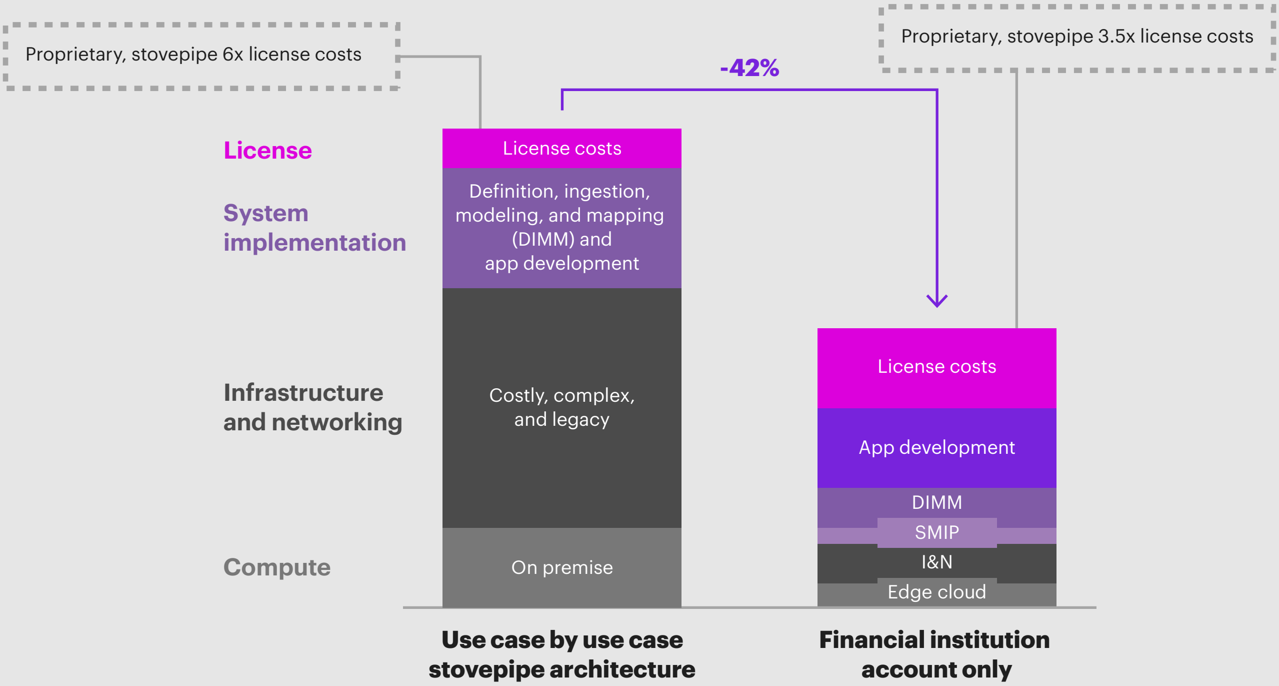

As artificial intelligence shifts from simple predictive models to autonomous, agentic AI, these SMIP capabilities are no longer required purely for integration. They are the essential infrastructure that enables AI to reason, collaborate, and deliver trusted decision-making based on provenance, lineage, and deep contextual integrity.

Sources: CESMII, ThinkIQ; Kearney analysis

The evidence: why most architectures fail to deliver scalable AI in manufacturing

Most manufacturers cannot scale AI because plant data lacks semantic consistency. As AI shifts from predictive models to autonomous agents, enterprises require a shared operational context—something that existing data architectures cannot provide. SMIPs dismantle this structural barrier and now sit at the center of AI-enabled manufacturing performance (see sidebar: Consumer foods company uses SMIP to unlock millions of dollars in economic value).

The cost of fragmentation and brittle integration

The current model forces companies into costly, repetitive integration work.

Cost and complexity. In an April 2025 survey of 500 business leaders, 74 percent said that infrastructure costs, disconnected data silos, or slow data ingestion were their biggest barriers to scaling AI, and 70 percent of organizations cited complex tool chains and fragmented data sources as their primary roadblocks.

Consumer foods company uses SMIP to unlock millions of dollars in economic value

The rapidly growing gluten-free food market, which is projected to reach about $13.7 billion by 2030, is highly sensitive to contamination risk. For manufacturers that process grains, trace gluten contamination can occur anywhere across a complex, multiparty agricultural supply chain, from crop rotation at farms to shared logistics and processing facilities.

A global producer of oat-based cereals faced this challenge in an environment spanning hundreds of assets and disconnected systems, where data existed but lacked shared meaning. Traditional analytics struggled with fragmented lineage and inconsistent data models, making root-cause analysis slow and unreliable.

By implementing a Smart Manufacturing Interoperability Platform, the company built a semantic backbone across the farm-to-fork value chain. Materials, batches, assets, and events were standardized and linked through a knowledge graph and material ledger, providing operational context and thereby, enabling full traceability of every grain movement and transformation. This contextual foundation allows AI agents to reason over material lineage, rapidly isolate contamination sources, and distinguish corrective actions from true root causes. Investigations that previously took weeks can now be completed in hours.

The impact was substantial: reduced recall risk, improved cleaning and yield efficiency, and tens of millions of dollars in economic value. More fundamentally, the case demonstrates that AI at scale requires shared operational meaning. SMIP provides the architectural layer that turns fragmented manufacturing data into trusted, actionable intelligence.

A lack of semantic consistency. Let's consider an example of a typical manufacturing plant without a standardized object information model. An electric motor in one production line might be named MOTOR_B01, while a similar motor in another production line is named MOTOR_AB. If you now want to deploy multiple applications (such as OEE tracking or predictive maintenance), you must manually customize every application for every motor, making cross-site learning and deployment impossible. A SMIP abstracts this problem by creating a consistent semantic foundation or language for the application layer. This interoperability is encapsulated with the release of CESMII's i3X API, which is an open, common API initiative to address the growing interoperability and scalability challenge in modern manufacturing architectures. The aim is for all SMIPs to be i3X-compliant so that vendor- and platform-agnostic applications can be built that use a common, standard API to facilitate interoperability at scale.

A lack of ontology and meaning. Manufacturing processes are complex with many inputs, pathways, and operational waypoints that influence the production of a final product. The data lineage of a batch or production unit is not provided by raw data. If it exists, it is often buried in siloed data systems. For agentic AI to derive meaningful and trustworthy insights from data, it needs to nderstand this provenance.

In most manufacturing environments, data is not structured to facilitate this. Data is often stored as flat tags, requiring manual interpretation with hard-coded analytics within site-specific silos. Figure 4 on page 6 and figure 5 on page 7 show these limitations based on a fictitious consumer packaged goods (CPG) manufacturing process for dishwashing liquid at Hand Dish Co.'s Berlin plant.

Historically, manufacturing data models have relied purely on hierarchical structures, the Purdue Model being the most influential (level 1 to level 4: control systems, supervisory systems, MES, and ERP). While this is effective for security and control isolation, it doesn't explain the nonlinear and interdependent nature of manufacturing processes and the complex interactions that describe information flow.

A SMIP provides this through an object-oriented information model, a knowledge graph, and material ledger that maps raw data to specific units of production, machines, materials, people, and events. Armed with this knowledge, agentic AI can truly understand the nuances of industrial manufacturing and deliver reliable, enterprise-wide performance improvement. Figure 6 on page 7 illustrates the application of a knowledge graph in two different contexts: a social media model and its analogous industrial model. Social media companies have been using knowledge graphs to drive insights and recommendations through a combination of subjects, relationships, and attributes. A SMIP enables this in a manufacturing context to provide end-to-end transparency on material flows, asset performance, and process outcomes.

We will now illustrate the value added by an object-oriented information model and a knowledge graph in our Hand Dish Co. example. Figure 7 on page 8 shows the information model for Hand Dish Co.'s Berlin plant. This is a structured, hierarchical model of manufacturing objects (such as assets, including sensors; materials; processes; and KPIs) that helps quickly scale manufacturing applications, analytics, and KPIs within and across sites. In this information model, asset hierarchy and process context is defined across four levels in the plant: enterprise and site, area, process cells or lines, and unit operations. For each unit operation (such as batch blending), the key process unit hierarchy and data context are detailed. They also provide the crucial plant-specific context for an AI LLM to provide actionable insights, grounded in live, plant data.

Figure 8 on page 9 illustrates a SMIP knowledge graph that models key relationships (such as asset can feed or operator-asset assignment). With a knowledge graph, an LLM will have complete genealogy of the process for any batch and can better monitor production and troubleshoot issues easily. While the information model in figure 6 provides a structured, hierarchical, and behavior-rich representation of Hand Dish Co.'s Berlin plant, the knowledge graph in figure 7 provides a flexible, relationship-centric way to connect and reason across entities.

Figure 4

Manufacturing data issues are preventing AI from scaling and delivering full value

Illustrative

Example: Overview of dishwashing liquid manufacturing process for Hand Dish Co.'s Berlin plant

Materials receipt and storage → Batch blending → In-process quality control and release → Packaging and palletization → Warehousing and shipping

Area: liquid blend

Equipment tag	Equipment or sensor	Parameters to measure
TK-001	De-ionized/soft water day tank	Level, conductivity
TK-101	Main mix tank with agitator/ load cells/jacket	Temperature, agitator speed, level, weight
SK-110	Surfactant tote/ IBC station	Flow
P-101	Transfer/ recirculation pumps	Status, speed
AIT-101	Inline pH analyzer	pH, temperature, calibration status
TT-101	Tank temperature	Temperature (PV, SP)
FT-102	Additive flowmeter	Flow rate
F-201	Bag/cartridge filter skid	Pressure differential
H-201	Finished product hold tank	Level, temperature

Area: PKG-L300

Equipment tag	Equipment	Parameters to measure
FL-301	Filler	Filling speed, temperature, pressure
CP-302	Capper	Torque, rejects
LBL-303	Labeller	Rejects
CW-304	Checkweigher	Weight stats, rejects
CS-305	Case packer and sealer	Rate, downtime
PLT-306	Palletiser	Rate, downtime

Source: Kearney analysis

Source: Kearney analysis

Figure 5

Four key issues with manufacturing data for an illustrative manufacturing process (Hand Dish Co, Berlin plant)

Issue #	Description	Issue details
1	Flat tags versus Object-oriented structure	—No inherent relationship between tags —No idea which area or unit operations an equipment tag belongs to —Impossible to generalise across plants
2	Manual interpretation versus Machine-understandable context	—Dataset from tank TK-101 may look like: —TT101 = 42.1°C —AG301 = 120 rpm —No clear indication if 42.1°C is good or bad —Which product is the dataset relating to? —Is the temperature and agitation speed quality critical?
3	Hard-coded analytics versus Reusable, scalable logic	—An Engineer may write a script for batch quality as follows: If TT101 > 50 and AG301 < 100: flag = "High Risk Batch" —Only works for tank TK-101 in Plant A —Breaks when name tags change and not reusable for tank TK-302, tank TK-705 etc
4	Site-specific silos versus Enterprise-wise consistency	—Plant A tag: TT101 = Pre-mix temp —Plant B tag: TMP101 = Pre-mix temp —Plant C tag: TANK3_TMP = Pre-mix temp —AI models trained at one site don't transfer to other sites seamlessly

Source: Kearney analysis

Figure 6

A knowledge graph helps agentic AI perform more reliably

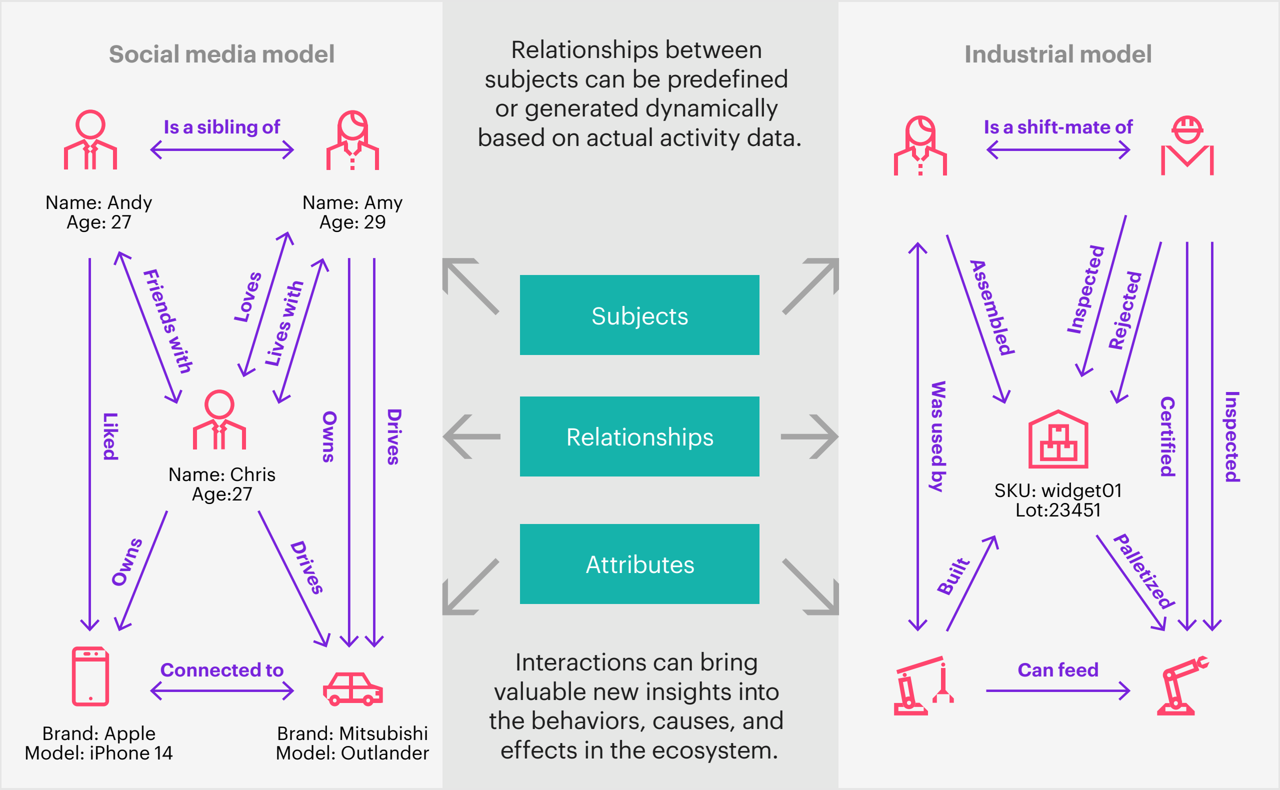

Source: Kearney analysis

Figure 7

Interoperability platforms allow raw data to be contextualised in an information model through process context, material flow and defined causal relationships

Example approach to object-oriented information model in consumer packaged goods example
(Hand Dish Co.'s Berlin plant)

Level	Detail	Example
1	**Enterprise and site**	Enterprise: Hand Dish Co Site: Hand_Dish_Co_Berlin
2	**Area** (to provide context: all assets relating to formulation)	CleanDishes1_Formulation
3	**Process cells and lines** (where batch genealogy, KPIs, digital twins anchor)	BrightShine_Spray_Line_1
4	**Unit operations** (class instance)	Unit 2: Batch_Blending Unit 4: Liquids_Packaging

Batch context

Batch	Activity	Value
BrightShine_Spray_2026_02_01	Processed In	**TK-001** DI/soft water day tank
	Processed In	**TK-101** Main mix tank with agitator/load cells/jacket
	Processed In	**SK-110** Surfactant tote/IBC station
	Processed Out	**H-201** Finished product hold tank
	Output product	BrightShine_DACH_12%

Process unit	Attribute	Value
Main_Mix_Tank_TK-101	Unit Type Max Volume Product Family	Agitated tank 10,000L BrightShine_DACH_12%
DI_Soft_Water_Day_Tank TK-001	...	...
Surfactant_tote_IBC_station_SK-110	...	...

Data context

Equipment (child object)	Property	Meaning/value
Temperature_sensor_TT101	CurrentValue EngineeringUnits HighLimit LowLimit Role CriticalToQuality	Live temperature °C 60 15 Process control Yes
Agitator_AG301	CurrentValue EngineeringUnits TargetUpperLimit CriticalToQuality	Rotation speed Rpm 120 Yes
Pump_P101	...	...

The information model is **semantic and hierarchical**:

- Process units and equipment/sensors have roles, properties, attributes, and limits defined.
- Process cells and unit operations know product families.
- Batches link everything together.

Source: Kearney analysis

Figure 8

A SMIP knowledge graph models key relationships

Illustrative

Example SMIP knowledge graph for Hand Dish Co.'s Berlin plant

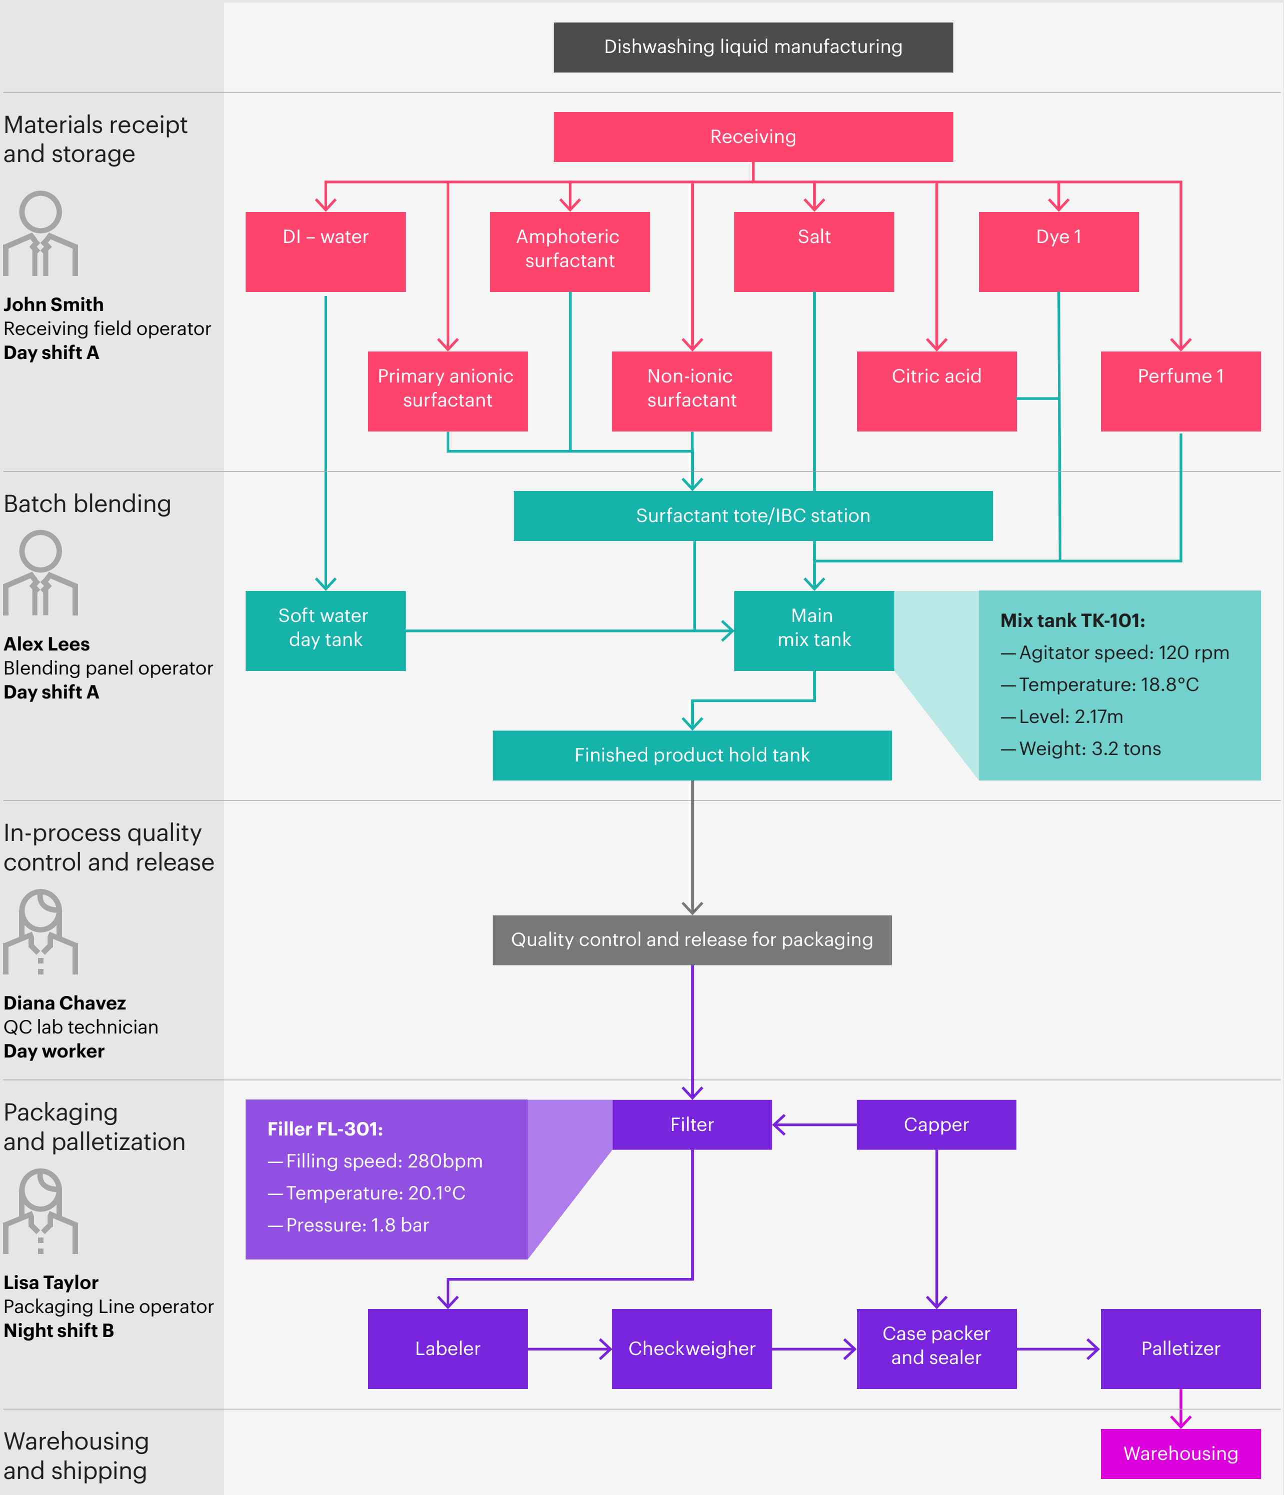

Note: SMIP is CESMII's Smart Manufacturing Interoperability Platform.

Source: Kearney analysis

When applied directly to raw data, AI can produce answers that sound right but are factually or operationally useless.

Next, we offer an example of two scenarios at our Hand Dish Co. plant on AI-assisted root cause analysis for a batch failure on a production line. Scenario 1 has an LLM on a data lake without an information model and a knowledge graph, while scenario 2 has this in place (see figure 9 on page 11).

"We've established that LLMs struggle with relational databases, especially when they ask for complex information across a diverse range of data sources. Integrating a knowledge graph gives LLMs a format they understand while broadening the depth and variety of data to pull from."

Forbes Tech Council, November 2024

Unreliable AI (hallucination risk). When applied directly to raw, non-contextualized data, LLMs and other AI can produce what is known as synthetic coherence: answers that sound right but are factually or operationally useless. For example, imagine an AI reviewing production data. Without context, it sees a correlation between a high defect rate and an operator reducing line speed. An AI that lacks semantic understanding might erroneously recommend running the line faster to reduce defects, confusing the operator's corrective action with the root cause. Contextual reasoning, made possible by a SMIP's knowledge graph, is the difference between shallow correlation and actionable causality. The risk of hallucination doesn't go away with larger models, more compute, and more data. Hallucination is a by-product of an LLM's freedom to "think for itself" rather than simply regurgitate information it has extracted from its training data. If we force an LLM to respond only to a prompt when it is 100 percent confident it isn't hallucinating, then it will refuse to respond to any prompts. The implication of this is that there will always be a non-zero hallucination risk with AI that needs to be managed. The scaffolding provided by a SMIP is one way of reducing the hallucination risk by forcing the AI to use the proper tools for navigating and extracting context and meaning from data.

Figure 9

Troubleshooting using AI at Hand Dish Co.'s Berlin plant

Illustrative

Scenario 1: LLM without an information model and knowledge graph

Packaging line operator detects overfilled and leaking bottles at the filler. Operator stops the line and asks the manufacturing chatbot to troubleshoot failure.

AI assesses available sensor data to generate an initial root cause analysis to guide troubleshooting. Operator has few potential immediate causes, but no verified cause.

Root-cause analysis highlights: Filler pressure set point and present value look normal. Last QC release (~45 min ago) shows all measured parameters are within limits.

Operator calls the process engineer and maintenance team to investigate upstream issues.

With limited interconnected data and relationships across equipment, the team must manually search through panel parameters to understand the root cause.

Outcome

- ✗ Operations team has limited understanding of root cause to prevent recurrence of issue (or effective troubleshooting in future).
- ✗ Batch frozen; materials likely scrapped (off-spec disposition).
- ✗ There is a long packaging equipment downtime for investigation and cleaning/line clearance prior to changeover to next batch.
- ✗ AI becomes another initiative that doesn't work.

Scenario 2: LLM with an information model and knowledge graph

Packaging line operator detects overfilled and leaking bottles at the filler. Operator stops the line and asks the manufacturing chatbot to troubleshoot failure.

Chatbot uses the genealogy of the specific batch with data from each preceding asset to trace dependent variables and potential causes of failure.

Chatbot hypothesizes that overfilling is potentially due to low liquid viscosity and high aeration from mix tank. It also flags high surfactant lot variability detected for current batch.

Chatbot recommends quality control viscosity check (as viscosity is not measured inline) to validate immediate cause.

Upon confirmation, chatbot recommends specific steps to remediate in batch blending to de-aerate liquid and resolve issue:

- —Reduce NaCl (salt) add rate by 10% in mix tank.
- —Reduce mix tank agitator speed to 100 rpm.
- —Increase de-aeration time to 20 min before quality control.
- —Monitor viscosity every 15 min; reset when normal.

Outcome

- ✓ Troubleshooting is faster and focused to isolate immediate cause and remediate root case. Operators not only resolve the problem faster but also understand it better for future reference and training.
- ✓ A smaller part of the batch is frozen (for affected time window); no need to scrap all materials. The downtime for production is shorter overall.
- ✓ AI enables troubleshooting and supports execution of corrective and preventative actions.

Note: LLM is large language model.

Source: Kearney analysis

Market implications: the strategic role of the SMIP

The core function of a SMIP is to provide semantic coherence across the entire operational landscape. It establishes a unified namespace authority: a shared, standardized taxonomy that defines what every asset is, its attributes, and its relationships.

Turning data into information: the context layer

A SMIP operates through a simple, three-tier model that acts as the Internet paradigm for manufacturing (see figure 10).

This aligns with CESMII's three smart manufacturing architecture imperatives for every SMIP (see figure 11 on page 13).

This abstraction of semantic consistency to the application layer is the foundational scaffolding for the coming age of agentic AI (see figure 12 on page 14).

The contextual bridge to agentic AI

Agentic AI represents the shift from passive predictive models to autonomous reasoning systems that can perceive, plan, and act. These agents (which might be assigned goals such as "optimize yield" or "reduce energy") cannot function without a common understanding of the operational world—a crucial ingredient that the SMIP provides.

The earlier example of a batch failure at Hand Dish Co.'s Berlin plant also illustrates this. Figure 13 on page 15 is a recap of the key elements of the problem.

Figure 14 on page 16 shows the various features of a SMIP, the requirement of agentic AI that it satisfies, and the impact it has on the quality and cost of analysis from the AI. We also illustrate each SMIP function for our example of a batch failure at the Hand Dish Co. plant described above.

Figure 10
Three-tier model

Tier	Layer	Detail
1	Application	Analytical and operational tools—such as digital performance management tools, AI models, root cause analysis, defect tracking, material ledger, carbon or cost racking apps—that consume information, not raw data
2	Context	Semantic, object-oriented model, and graph model that gives meaning to the data, describing what each signal represents, how it relates to other entities, and how it fits into the overall process
3	Data	Raw signals, sensor readings, machine tags, and system events generated at the edge

Source: Kearney analysis

Figure 11
CESMII’s three imperatives for a SMIP

Imperative	Detail
Information model standardization (smart manufacturing profiles)	Open, standards-based information-modeling strategy for manufacturing (and related supply chain) devices, assets, and processes
Contextual manufacturing information platform	A clear set of requirements enabling manufacturing platform and application interoperability
i3X	An open, common API for manufacturing systems, rapid app development, scaling AI deployments, enterprise application integration, and supply chain optimization

Note: SMIP is CESMII’s Smart Manufacturing Interoperability Platform.
Source: Kearney analysis

Figure 12
Three-tier model with typical elements in each layer

Application

Agentic AI and application layer

- Agentic AI (LLM query database and model)
- Agentic AI (root cause analysis)
- Agentic AI (autonomous workflow)
- Real-time analytics (AI-driven decision-making)
- Eternal MCP servers (third-party tools)
- Multi-agent workflows (task-driven execution)
- Applications (ABC, quality, carbon tracking etc.)

Context

Model context protocol (MCP): server and clients

APIs

Semantic framework and ontology

- Object-oriented information model (type definitions)
- Knowledge graph model (relationships)
- Nextgen historian (postgres database)
- Attribute tag mapping (abstraction layer)
- Cloud libraries (SM profiles)
- Unified namespace
- Material ledger (data provenance)

Interoperability and standardized connectivity layer (MQTT, OPC-UA, Custom Connector, etc.)

Data

Existing legacy OT	Shop-floor controls devices	Manufacturing systems	Enterprise/ management systems
activplant FactoryTalk Rockwell Automation VTScada	PLC	Local historian MES/CMMS/QMS etc. HMI/SCADA IIOT/sensors	ERP WMS/IM Regulatory PLM etc.

Source: Kearney analysis

Figure 13

Details of batch failure on packaging line in Hand Dish Co.'s Berlin plant

Symptom	During a packaging production batch run (Batch # FG-DW-ULTRA-ORIG-473ML), the operator observed the following: — Leak tester/vision rejects spike (cap leaks/weepers) from ~0.2% to 6% — Checkweigher variance increase (more overfills) Filler operating parameters looked normal and capper torque setpoints unchanged; mechanical checks didn't find a single failed head or nozzle
Causes of the problem	The bulk liquid batch (Batch # BULKLOT-26-0302-001) feeding the packaging production batch had low viscosity due to high entrained air (foam) caused by a combination of: — **Surfactant lot variability** (different thickening response vs salt curve) — **NaCl (salt) addition executed too fast** (addition rate deviation) — **Mix tank agitator RPM scaling changed** after recent maintenance, increasing shear/ aeration in liquid — **De-aeration/hold time shortened** to keep packaging line supplied continuously
Impact of the problem	Foamy product was causing inconsistent fills, product on neck finish and cap threads. Seal integrity failures even though packaging equipment looked "in control."

Source: Kearney analysis

Figure 14

A SMIP-enabled AI operating system is the scaffolding for agentic AI in manufacturing

SMIP function	Agentic AI Requirement	Impact	Example from root-cause analysis of a batch failure at Hand Dish Co, Berlin plant
Semantic modeling: ontologies, object models, and knowledge graphs	Common world model (digital twin)	Enables agents to reason about entities (e.g., equipment, material, and batches), not just flat data	Agent can move instantly from packaging line (PKG-L300) rejects → current work order # → finished goods lot(s) impacted → bulk lot feeding filler → originating batch (batch-...) and pulls the relevant segment executions (“Add NaCl”, “De-aerate/ hold”, “Transfer”). No tag-hunting is involved. The agent knows the meaning of each signal in the process.
Federated information exchange	Local autonomy, global coordination	No need to put all relevant data into the cloud. Agents can be deployed at the data source, massively reducing cloud storage costs.	Packaging agent queries in place (at the edge) for: filler trends, capper events, leak tester classifications; simultaneously queries blending cell/historian for agitator RPM, pump states, salt addition rate, and queries receiving/ERP for surfactant lot used, without moving full historian datasets to the cloud. Only “answers” (context + features) are exchanged.
APIs for real-time context access	Real-time situational and context awareness	Agents can “see” current plant state through contextual queries	Agent executes contextual queries such as: “Show last 30 min of CW-304 weight distribution, FL-301 nozzle status, and the bulk feed context (H-201 level, recirc status, bulk temp, latest lab viscosity) for the active BULK lot.” It detects that the increase in rejects correlates with a step change in bulk handling (e.g., recirculation on / higher shear) and not a packaging head failure.
Provenance and versioning	Trust and explainability	Agents have access to the full data lineage as material moves through the supply chain.	Agent validates which recipe version ran, whether the mix tank agitator scaling changed after maintenance, and whether the salt feeder scale or pH probe had calibration flags. It explains the conclusion with auditable evidence: “NaCl add-rate exceeded recipe limit by X percent during SE-005; agitator RPM mapping changed on date/time; de-aeration segment duration reduced vs standard.” This prevents a “blame the filler” misdiagnosis.
Relationship graphing	Multi-agent collaboration, cause and effect analysis	Supports causal reasoning (“if this, then that”) across processes	Agent correlates across domains: leaker events ↔ bulk viscosity + deaeration duration ↔ mix tank agitator RPM/shear ↔ surfactant lot genealogy ↔ maintenance change record. The relationship graph highlights: “Rejects occur only on BULK lots containing surfactant SLES-LOT-X AND batches run after mix tank agitator RPM scaling update.” It recommends targeted corrective actions: “Enforce minimum de-aeration time, lower NaCl add-rate, restore agitator scaling, and (optionally) adjust recipe salt curve by surfactant SLES lot attributes.”

Note: SMIP is CESMII’s Smart Manufacturing Interoperability Platform.

Source: Kearney analysis

Strategic utlook: executive decisions for the age of SMIPs

The shift from a data-centric to context-centric architecture is a strategic imperative that directly impacts cost, risk and competitive advantage.

Architectural migration: when to move

The big data paradigm assumed that massive data lakes would lead to insight. The reality is that manufacturing data without context is next to useless.

When to migrate. If your enterprise is struggling to move AI pilots into production, if digital transformation projects are consistently delayed by integration rework, or if your tribal knowledge of plant data is concentrated in a few departing data scientists or engineers, these are compelling imperatives to invest now in interoperability and standardization.

The shift. Stop moving petabytes of raw data to the cloud for slow, expensive analysis. Instead, deploy a SMIP to move meaning: lightweight, structured descriptions of the data across a distributed ecosystem. This transition changes the cost structure of analytics by moving from manual data wrangling to autonomous, agent-driven insight.

Sequencing and standardization

The SMIP is the platform for industrial digital scalability, allowing best practices and intelligence to be replicated rapidly across sites.

Where to standardize first. Focus semantic standardization on assets and processes that are common across multiple plants, such as motors, pumps, CNC machines, and reactors, and those that are critical to strategic manufacturing goals, such as yield optimization, throughput increase, energy consumption reduction, improved carbon tracking, and activity-based costing). Use object-oriented modelling to define digital objects that inherit attributes and behaviors, making the model reusable across vendors and plants.

The governance model. The SMIP's unified namespace (UNS) becomes the single source of truth for nomenclature and taxonomy within the enterprise. For example, a level on a tank is the level attribute on the respective tank object in the information model rather than a LI21324.pv. This abstraction allows the applications in the application layer to focus on function and impact rather than integration complexity.

Competitive advantage and risk

Early adopters of a SMIP can gain a decisive advantage, while laggards incur significant risk in fully leveraging AI in manufacturing. Figure 15 on page 18 shows these key areas of advantage and risk.

The logical end game of this transformation is cognitive manufacturing: factories that can think, learn, and adapt autonomously. For the next three to five years, the competitive edge will belong to enterprises that successfully transition from data-centric to context-centric operations thus unleashing the untapped potential from agentic AI. Without the ontological and semantic framework that a SMIP provides, agentic-AI deployments will be not only struggle to scale but will be unable to extract meaningful insights for enterprise value.

Figure 15
Competitive advantage and risks

Early adopter advantage	Laggard risk
Accelerated AI deployment: AI models are instantly context-aware and scalable.	Architectural debt: perpetual, bespoke integration rework and unscalable pilot projects
Autonomous optimization: agents deliver real-time, cross-site learning and efficiency.	Inconsistent governance: inability to compare performance across sites due to disparate data definitions
Reduced risk: semantic grounding eliminates AI hallucination and ensures traceability and auditability.	Operational blindness: insights that are too slow, too expensive, and often too late to act upon

Source: Kearney analysis

Authors

Sumeet Ladsaongikar
Partner, London
sumeet.ladsaongikar@kearney.com

Parshu Sankaranarayanan
Consultant, London
parasuraman.sankaranarayanan@kearney.com

Paul Turner
Chief Technology Officer, ThinkIQ

For 100 years, Kearney has been a leading management consulting firm and trusted partner to three-quarters of the Fortune Global 500 and governments around the world. With a presence across more than 40 countries, our people make us who we are. We work impact first, tackling your toughest challenges with original thinking and a commitment to making change happen together. By your side, we deliver—value, results, impact.

kearney.com

KEARNEY

100 Years of **Impact**