From manufacturing execution systems to *intelligence*

Powering AI-driven pharma operations

Why the next value horizon in pharmaceutical manufacturing depends less on the models you choose than on the manufacturing context you can give them.

Make it real.

Capgemini

Executive summary

Pharmaceutical manufacturers have spent a decade making execution digital. The constraint on AI is no longer data capture. It is meaning.

In the most mature networks, the batch record is electronic, the process is enforced, the audit trail is complete. And yet when a deviation is raised at two o'clock in the morning, an investigator still spends the first part of the investigation doing something no one designed: assembling evidence by hand from multiple systems that do not share a definition of the batch.

That gap is the subject of this point of view. Manufacturing data carries almost no shared meaning outside the system that produced it, and, absent that meaning, AI cannot be trusted to act. Industry benchmarking now identifies data readiness as the single largest barrier to AI deployment across the sector, with more than half of surveyed organizations yet to harmonize their unstructured data.[1].

The results follow the constraint. In a mid-2026 survey of 150 life sciences executives, 45% reported that AI initiatives had produced measurable improvement, but only 13% reported measurable improvement at scale.[2] In parallel benchmarking, AI was live and embedded in only four of 30 recorded manufacturing and quality use cases.[1] The pilots work, but they do not become operations.

This is precisely where manufacturing execution systems (MES) matter. Manufacturing intent—this batch, this recipe step, this material lot, this equipment train, this operator action, under this approved procedure—is what MES holds and almost nothing else does. MES is not a legacy system to be worked around on the way to AI; it is the semantic and procedural foundation that makes agentic AI defensible in a GxP environment.

The value is not secured in the transaction itself. It is won or lost in the transition that follows. Value is won or lost based on how well leaders plan the operating model, data, testing and non-functional requirements before execution begins.

"AI models lack a true semantic understanding of manufacturing intent." BioPhorum industry benchmarking, BioProcess International, July 2026.

This point of view demonstrates how organizations can:

- Use MES as the trusted source of batch, process, material, equipment, and quality context.
- Connect operational data through semantic, reusable, and validated information foundations.
- Apply AI agents to assemble evidence, recommend actions, and coordinate workflows inside the flow of work.
- Preserve human accountability, explainability, and GxP controls.
- Scale from focused use cases to enterprise capability through common architecture and governance.

The immediate value sits in problems every manufacturing and quality leader recognizes: deviation triage, corrective and preventive action (CAPA) orchestration, batch release readiness and frontline decision support. The longer-term prize is a connected manufacturing network in which people and intelligent agents work together to improve right-first-time performance, release speed, productivity, and resilience.

One thing has changed since we last published. In January 2026 the FDA and EMA jointly issued ten guiding principles for good AI practice across the drug product life cycle, explicitly including manufacturing[3]. With these directives, governance is no longer the reason to wait. It is the specification to build against.

The value at stake

Reported outcomes when MES is executed alongside process and technology change:

Metric	Reported outcome
Manufacturing lead time	50% reduction
Overall equipment effectiveness (OEE)	25–45% gains
Documentation time	75% reduction
Right-first-time performance	35% improvement
Quality-related costs	25–30% drop
Documentation errors	95% fewer

Source: *European Journal of Computer Science and Information Technology, 2025[4]. Figures are reported industry outcomes where MES is deployed in conjunction with wider technology and process improvement; they are not a forecast of any individual organization's result.*

Introduction: From MES value realization to intelligent operations

Pharmaceutical manufacturers have invested significantly in manufacturing execution systems to digitize production, strengthen compliance and standardize execution. The next opportunity is to extend these investments beyond transactional control and use MES as the foundation for intelligent operations.

This changes the question leaders should be asking, from "is our MES deployed?" to "what decisions does our MES now make possible that were not before?" Value realization is measured in outcomes—right-first-time performance, deviation reduction, batch-release readiness, operator productivity, and resilience—not in functionality delivered at go-live.

Much of the industry's digital manufacturing journey has focused on converting manual processes into electronic ones. Electronic batch records, digital workflows, and automated documentation have created meaningful benefits in compliance, traceability, and data quality. Yet many of the fundamental operational challenges that existed before digitization remain. Production delays, recurring deviations, lengthy investigations, and inconsistent performance improvement efforts are often no less common simply because the paperwork became electronic.

This reflects an important shift in where value can now be created. For many pharmaceutical manufacturers, the constraint is no longer access to data. MES, historians, laboratory information management systems (LIMS), quality management systems (QMS), enterprise resource planning (ERP) systems, and connected equipment already generate more information than most organizations can effectively utilize. The challenge now is not capturing information but rather applying it.

The next phase of manufacturing transformation therefore is not primarily about digitization. It is about turning operational data into better decisions, faster learning, and measurable improvement. That transition from information to action is where intelligent operations begin.

For more than a decade, pharmaceutical manufacturers have invested heavily in MES and related operational technologies. Those investments created trusted execution records, automated large portions of manual activity, improved compliance, and generated enormous amounts of manufacturing data. In many respects, the industry has already built the foundation it needs.

The opportunity is to capitalize on that investment. Advances in connectivity, contextualization, semantic architectures, and agentic workflows make it possible to use manufacturing information in ways that were previously impractical. The objective is no longer simply to record what happened, but to apply what has been learned to improve performance, solve problems more effectively, and support better operational decisions.

This is why the discussion around AI in manufacturing is so important. The industry has the data. The question is whether manufacturers can finally transform that information into measurable operational improvement.

Why AI has not delivered its full potential

The principal constraint is not the availability of algorithms but the absence of an integrated operating foundation. Manufacturing data remains fragmented across automation systems, historians, MES, LIMS, QMS, ERP and unstructured records, and the context needed for reliable decisions is often incomplete or inconsistent.

Benchmarking reaches the same conclusion from the other direction: of 19 recently implemented AI tools in this space, 12 were built entirely in-house and none purchased off the shelf, a signal that the hard part is local context and integration, not the model[1].

The consequence is a recognizable failure mode. Insight is generated outside the flow of work, in a tool a decision-maker does not use, expressed in terms a quality reviewer cannot accept as evidence. It cannot translate into trusted, repeatable, governed action, so it does not scale, and over time it is not maintained.

What this point of view will demonstrate

How MES evolves from a system of execution into a strategic enabler of AI-driven operations; how a Unified Namespace (UNS) and enterprise data platforms extend interoperability; and how agentic capabilities detect events, analyze conditions, recommend responses, and orchestrate workflows within defined GxP guardrails.

Throughout this point of view, anonymized examples from Capgemini show what has actually been built, at what scale, and what it took. The argument is deliberately practical. Every claim is tied either to published evidence or to work delivered in a live GxP environment, and every recommendation is one a manufacturing or quality leader could act on in the next planning cycle.

Regulatory context: A shared baseline for AI in manufacturing

On January 14, 2026 the FDA and EMA jointly published Guiding Principles of Good AI Practice in Drug Development, which contained ten principles covering AI used across the drug product life cycle, explicitly including manufacturing.[3]

When read as an engineering specification rather than policy, **five are immediately actionable**:

Risk-based approach	validation proportionate to model risk.
Clear context of use	one named decision, one named owner.
Data governance	provenance and processing traceable.
Performance assessment	evaluate the whole system, including human–AI interaction.
Life cycle management	monitoring and re-evaluation for drift.

MES as the foundation for intelligent operations

Why manufacturing context and governed execution matter

AI becomes operationally useful only when data is connected to the correct product, batch, recipe, material, equipment, process step, operator action, and quality status. A temperature excursion is a number; a temperature excursion during the hold step of a specific batch, on equipment with an open maintenance work order, is a decision.

MES supplies that difference. It translates approved procedures into controlled workflows through enforced sequencing, parameter checks, electronic records, exception handling and approvals. This is what allows an AI recommendation to be traceable, procedure-aware, and subject to clear human accountability. It is the difference between a suggestion and admissible evidence.

More than a system of execution

A mature MES standardizes processes, embeds quality controls, creates a trusted digital production record and connects automation, historians, LIMS, QMS, ERP and data platforms. Its value should be measured through outcomes, not functionality.

The mechanics are concrete. Review by exception is the clearest illustration: the shift from reviewing a complete batch record page by page to reviewing a short exception report with the material items highlighted is what moves batch review cycle time from days to hours[5]. Analysis of digitized quality operations has associated this class of change with reductions in overall deviations of more than 65%, closure times more than 90% faster, and the elimination of up to 80% of manual documentation work[6].

Data, intelligence and action

MES need not contain every data set or AI model and attempts to make it do so are a reliable source of cost. Its distinctive role is narrower and more valuable: to supply process truth, execution status, and controlled interaction points.

Each use case should establish a baseline, a target improvement, and a named benefit owner before deployment. Use cases that cannot name all three are not ready.

Value should be tracked through a small, stable set of measures: batch-release cycle time, right-first-time performance, deviation frequency and closure time, review effort, operator productivity, and repeat-event reduction.

Case Study

Scaling MES across more than 50 sites

Standardization as the engine of value

A top-five global pharmaceutical manufacturer executed a MES and electronic batch record transformation spanning more than 50 sites and 25 production processes.

Delivery used a centralized “MES factory” model to industrialize configuration, deployment and validation, supported by machine-connectivity audits, SOP updates, operator training, and a harmonized governance structure. Standardized templates and a global delivery model enabled rapid rollout while still allowing site-specific configuration.

Why it matters

The reusable core (templates, governance, validation patterns), is what compresses timelines. The individual site project never does.

Five stages of MES maturity

Value compounds across five stages. They are cumulative, meaning that organizations attempting stage four without stage two consistently rediscover that harmonization, governance, and adoption are not optional prerequisites but the mechanism by which value is retained.

Stage	Focus	What it makes possible
1.Stabilize	Compliant, reliable execution; downstream incident and problem management	A trustworthy, credible production record
2.Harmonize	Standardize processes, global core model, adoption and competence	Comparability across sites — the precondition for any reusable agent
3.Integrate	Connection to ERP, LIMS, QMS, historians, CMMS; performance optimization	Cross-system context: an event can be understood, not just recorded
4.Predict	Predictive and assistive intelligence embedded in the workflow	In-workflow guidance rather than retrospective reporting
5.Govern & orchestrate	Coordinated agentic operations across production, quality, and engineering	Governed evidence, decisions and action at network scale

Capgemini MES maturity model. Stage descriptions reflect delivery experience across global pharmaceutical manufacturing networks.

How to use the model

Measure the stage you are in, not the one you aspire to.

A short, stable set of outcome measures: track batch-release cycle time, right-first-time performance, deviation frequency and closure time, review effort, operator productivity, and repeat-event reduction. Each intelligent-operations use case should establish a baseline, a target, and a named benefit owner before deployment.

Stages are cumulative.

Attempting predictive or agentic capability on top of unharmonized processes produces results that cannot be reproduced at the next site, which is the most common and most expensive way to discover that stage two was skipped.

Why it matters.

Maturity is a diagnosis, not a target. It tells you which constraint to remove next.

Case Study

Deciding what a batch record should contain

Sterile manufacturing, global vaccine manufacturer

Approximately 15,000 batch-record data points across sterile manufacturing were analyzed to separate critical from non-critical data before transposition into MES. The work spanned antigen, purification, filling, inspection, solutions preparation, and packaging.

The team mapped the batch review process and its KPIs, clarified the role of content across the record, defined the target data type for each parameter, and validated a block structure and recipe architecture with production and quality before configuration began.

Why it matters

Lean batch record design is a data-governance decision taken before the MES is configured. Taken afterwards, it is a remediation program.

Why AI initiatives struggle to scale

The barriers are consistent and, by now, well documented. Sustainable scale requires organizations to treat AI as a governed business capability rather than a portfolio of experiments.

Data fragmentation and context gaps

Relevant data is distributed across MES, LIMS, QMS, historian, ERP, instruments and documents, often with inconsistent terminology and limited lineage. Without common definitions for products, batches, equipment, process steps, samples and quality events, models remain site-specific and difficult to reuse, which is why so many succeed once and never again.

This is now the sector's binding constraint rather than one issue among several. A governed semantic layer and reusable data products convert fragmented information into decision-ready context that travels across systems and sites.

It is worth being precise about where the effort goes. One widely used framing puts 10% of the challenge in the algorithms, 20% in data and technology, and 70% in people and processes[7]. Manufacturing AI programs that invert that ratio in their budgets tend also to invert it in their results.

Governance, validation and trust

In GxP operations, model accuracy alone is insufficient. Organizations must define the intended use, data controls, acceptance criteria, explainability, audit trails, human oversight, change management, and ongoing performance monitoring.

The January 2026 FDA–EMA principles make this expectation explicit, and in doing so remove a common excuse: performance assessment is expected to evaluate the complete system including human–AI interaction, and quality management is expected to run across the AI life cycle with scheduled monitoring for drift[3]. Retrofitting either onto a finished pilot is materially more expensive than designing for both.

Moving beyond isolated pilots

A pilot proves that something can work; industrialization proves that it can work again without the original team. Most pharmaceutical AI portfolios are rich in the first and thin in the second.

Scaling begins with a high-value operational decision, a named business owner and measurable outcomes. The solution must be embedded in existing workflows, supported by reusable architecture, governed data products, and clear adoption responsibilities. A focused lighthouse use case can prove value, but industrialization requires standard templates, lifecycle governance, monitoring, funding, and a replication roadmap.

Case Study

Driving decision intelligence for a global pharmaceutical leader

Formulation manufacturing yields varied significantly, from 60% to 80%, while production, quality and maintenance data remained fragmented across PAS-X, SAP, LIMS, historians, IoT/OPC-UA and laboratory systems. This limited visibility into root causes and constrained performance improvement.

The program established a unified manufacturing data model and ontology covering batches, equipment, materials, processes and quality. Data from operational and enterprise systems was integrated into an AI-powered manufacturing intelligence platform, with an industrial knowledge graph connecting assets, process parameters, events, and quality information. Analytics and AI-enabled applications supported yield optimization, root-cause analysis, asset performance, and quality monitoring, using a scalable blueprint designed for deployment across global sites.

Reported outcomes included a 5–10% increase in yield, 20–30% faster root-cause investigations, and a 10–20% improvement in asset effectiveness, alongside improved batch consistency, compliance, and audit readiness.

Why it matters

Combining contextualized industrial data, knowledge-graph technology, and AI-driven analytics created faster decisions, measurable operational improvement, and a repeatable foundation for enterprise-wide manufacturing intelligence.

Where value is landing today

In regulated operations, AI adoption is currently concentrated in a small number of high-value workflows, including deviation management, documentation review, compliance checking and complaints management. The greatest gains come from accelerating evidence gathering, document review and information synthesis, while keeping human judgement firmly in the loop.[1]

16 hrs → <1 hr

Reduction in MES source-code review time achieved in a documented implementation.[1]

From insight to action: Agentic AI

Traditional analytics explains what happened or predicts what may happen. Agentic AI interprets context, assembles evidence, recommends actions, and coordinates governed workflows.

It creates an action layer over MES and connected manufacturing systems, enabling faster problem solving while retaining human accountability for GxP-impacting decisions. Realism is warranted about how far this has travelled: industry consensus is that fully autonomous systems managing complex regulated workflows remain some way off, and human-in-the-loop oversight remains essential[1]. That is not an argument for waiting, but for building agents whose value does not depend on autonomy—which is where most of the value is.

What makes agentic AI different

An agent works toward a defined operational objective rather than producing a single response. Unlike conventional automation, it can adapt its investigation path to the situation. Unlike a standalone copilot, it can coordinate multiple steps across a workflow. Its actions remain bounded by approved tools, explicit decision rights, audit trails, and human approvals.

MES enabling the agentic journey

MES provides the batch, recipe, process-step, material, equipment, exception, and user context that allows agents to reason about manufacturing events accurately. An agent does not need to infer manufacturing intent if a governed system already holds it.

Connected with QMS, LIMS, historians, CMMS and knowledge repositories, MES lets an agent move from detecting an event to understanding its operational significance; then it can assemble evidence, propose the next best action, and initiate a controlled workflow without duplicating records or bypassing established quality processes.

From analytics to intelligent operations

Descriptive	shows what happened.
Predictive	indicates what is likely to happen.
Prescriptive	recommends what should be done.
Agentic	coordinates evidence, decisions, and approved actions across systems and roles.

Each stage changes who acts on the output, and that is the real transition. Descriptive analytics produces a report for an analyst; agentic operations produce a task for an accountable role, with the evidence already attached. The objective is not unrestricted autonomy but faster and more consistent decision execution within defined procedures, controls, and approval boundaries.

Accelerator

The pattern: manufacturing problem solving agent

Deviation triage and CAPA orchestration

When an issue is detected, the agent uses MES batch and process context, historian trends, laboratory results, equipment history, operator comments, SOPs and previous deviation records to create an evidence package.

It classifies the issue, identifies missing information, retrieves similar cases, proposes ranked root-cause hypotheses, and drafts investigation or CAPA tasks for review.

It does not approve conclusions or close GxP records. Investigators review the evidence, accept or revise the hypothesis, and authorize the next action.

Built as a pattern rather than a tool, the same design extends to batch-release readiness, troubleshooting, and throughput constraints.

Building the intelligent manufacturing stack

Intelligent operations require a layered technical core rather than a single platform. No layer should be asked to do another layer's job.

MES governs execution and records process context; a Unified Namespace enables event-driven exchange; data platforms provide scalable storage, contextualization and analytics; and AI converts trusted information into insight and governed action. Most of the architectural debt in manufacturing digitization comes from ignoring that separation under schedule pressure.

The role of Unified Namespace

A UNS is a semantic, event-driven information layer—not a replacement for MES, automation or enterprise systems. It allows programmable logic controllers (PLC), distributed control systems (DCS), supervisory control and data acquisition (SCADA) systems, historian, MES, LIMS, QMS and ERP to publish and consume consistently named, contextualized events.

Shared semantics, security, lineage, and validation connect these capabilities without displacing system ownership or transactional control.

The benefit is structural rather than immediate. Point-to-point integration cost grows roughly with the square of the number of endpoints; a publish-and-subscribe semantic layer grows roughly linearly. On a single site the difference is arguable. Across a network of plants with differing automation vintages, it decides whether use cases can be replicated at all.

The role of data platforms

Data platforms retain historical and cross-functional information, harmonize data, manage lineage and support analytics, knowledge graphs, digital twins, and AI model lifecycles. The UNS answers "what is happening now, and what does it mean?" The data platform answers "what has happened before, and what does that imply?" Agentic reasoning needs both.

Closing the loop

An event is generated on the shop floor, enriched with MES and semantic context, analyzed against history, and evaluated by AI. The recommendation is presented with evidence and routed through an authorized workflow. The loop is the deliverable; an architecture that produces excellent recommendations but has no governed path to action is a reporting system with higher running costs.

Approved actions are executed through the appropriate operational system, and the outcome is captured to improve future decisions. That last step is the one most often skipped, and it is what turns a deployment into a capability that compounds.

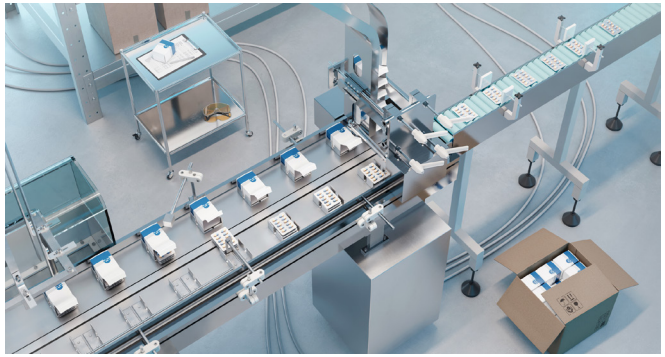

Case Study

Preserving meaning from edge to cloud

13 plants, global pharmaceutical manufacturer

A heterogeneous environment across 13 drug-substance and drug-product plants, several acquired, with widely differing digital maturity, required a scalable, GMP-ready manufacturing data foundation, compliant from day one to avoid rework.

The architecture combined an edge unified data exchange layer with an enterprise semantic platform, using shared models based on manufacturing and laboratory standards.

The initial four-month implementation covered two plants and delivered schedule-adherence and production-insight use cases on a validation-by-design platform. Reusable semantics then accelerated deployment and a further use case linking batch and time-series data was delivered within the same week.

Why it matters

A data foundation earns its cost by preserving meaning, not by centralizing data.

One authority per layer

This separation of concerns lets each component retain its authority while creating a scalable foundation for intelligent applications.

Layer	Authority	Explicitly not responsible for
OT and instruments	Generating events at source	Meaning, process context or retention
MES	Governed execution—batch, recipe, material, equipment, exception, approval	Being the data lake or the model host
Unified Namespace	Distributing semantic, real-time information across operations	Long-term history or heavy computation
Data platforms	Retaining and enriching historical and cross-functional data; model lifecycle	Real-time execution control
AI agents	Interpreting combined context, presenting evidence, coordinating approved action	Approving GxP-impacting decisions
Cross-cutting controls	Security, identity, lineage, validation, monitoring and audit across every layer	–

Case Study

An event-driven backbone across a network

20+ sites in 15 countries

A bi-directional data flow service connects shop floor and cloud across more than 20 manufacturing sites in 15 countries. Data from industrial equipment is captured, normalized, contextualized, and enriched at the edge, then integrated with enterprise applications.

The topology pairs on-premises brokers at each site with a centralized cloud broker. Architecture principles were set up front to be event-driven, real-time, secure and scalable.

Why it matters

A per-site broker feeding a central hub lets a network add a site without renegotiating its integration model. That is the difference between a coordinated rollout and 20 projects

Operating agentic AI in GxP environments

The objective is controlled augmentation: agents collect evidence, apply approved rules, explain recommendations, and coordinate tasks, while qualified people retain accountability for regulated decisions.

Human oversight and accountability

Decision rights should be explicit for every use case. Agents may draft, recommend, and initiate approved workflows, but authorized users must review product-impact conclusions, severity classification, root cause, CAPA plans, effectiveness checks, and batch disposition.

The design test is simple and worth applying early: for every output the agent produces, name the role that owns the resulting decision. Outputs with no owner should not be produced.

Explainability and auditability

Each recommendation should show the evidence used, the applicable rule or rationale, confidence, missing information, and assumptions.

The audit trail should retain source identifiers, agent outputs, human changes, approvals, workflow transitions, and timestamps.

The standard to design to is reconstruction: A reviewer must be able to establish not only what was decided but how, on what basis, and by whom—months later, without access to the team that built the agent. An agent that cannot show its evidence has not saved an investigator any work, because the evidence will need to be gathered again.

Governed decision making

Controls should be proportionate to the risk of the decision. Role-based access, least-privilege tools, validated interfaces, approved knowledge sources, model monitoring, and defined stop conditions prevent silent changes and unsupported conclusions.

One rule does more work than any other: deterministic logic must remain the authority for classification matrices, authorization checks, and workflow states. AI supports evidence retrieval, explanation, and drafting. Where a decision must be reproducible on demand, it should not be probabilistic—a design choice that also makes validation dramatically cheaper.

Controls of this kind are not friction. They are what allows a manufacturing organization to put an agent in front of a regulator and explain, line by line, how a conclusion was reached.

Proof of Concept

Agent-assisted deviation triage

Global biopharmaceutical manufacturer (proof of concept)

The agent supports a deviation from initial triage through controlled CAPA planning. It assembles evidence from QMS, MES, historian, LIMS, CMMS and controlled documents; suggests severity with visible rationale; retrieves similar cases; drafts investigation content; and proposes CAPA tasks, owners, evidence needs, and effectiveness checks.

QA and subject-matter experts review every regulated conclusion, and the system records the complete decision trail. No GxP record is closed autonomously

Why it matters

The constraint that no record closes autonomously is not a limitation on the value. The evidence-assembly phase is where the time is lost, and that phase needs no autonomy at all.

Mapping the FDA–EMA guiding principles

Ten principles, published jointly in January 2026, read as an engineering specification for manufacturing agent design.

Principle	What it means for a manufacturing agent
Human-centric by design	Agents draft and recommend; qualified people decide. Every output has a named accountable role.
Risk-based approach	Control envelope and validation effort scale with the risk of the decision, not the sophistication of the model.
Adherence to standards	GxP, cybersecurity and data integrity are design inputs, not a downstream review.
Clear context of use	One defined decision, one defined scope, documented before build.
Multidisciplinary expertise	Manufacturing, Quality, IT, Data and Engineering own the agent jointly.
Data governance and documentation	Provenance and processing steps traceable; the semantic layer is how this scales.
Model design and development	Fit-for-use data; explainability weighted alongside predictive performance.
Risk-based performance assessment	Evaluate the whole system, including human–AI interaction.
Life cycle management	Scheduled monitoring, periodic re-evaluation and drift detection.
Clear, essential information	Plain-language context of use, performance and limitations for the people who use it.

Principles as published in FDA and EMA, Guiding Principles of Good AI Practice in Drug Development, January 2026.[3] The right-hand column is Capgemini's interpretation for manufacturing agent design, not regulatory guidance.

Case Study

Turning multivariate statistics into an operator alert

Global pharmaceutical manufacturer

Multiple process parameters influencing product quality were logged into the process historian and reduced, through principal component analysis, to two components that characterize the behavior of the process. An acceptable operating region was defined from historical batches that produced good quality outcomes.

Real-time monitoring evaluates the live parameter combination against that region. A deviating combination is flagged as it emerges, and a notification is raised so the operator can take corrective action within the batch rather than discovering the problem at review.

Why it matters

This is prediction delivered as an in-workflow instruction rather than a dashboard—and it works without any autonomy whatsoever.

Manufacturing problem solving in practice

Many manufacturing investigations exist primarily to satisfy compliance requirements. As a result, teams can become highly effective at documenting events, assigning actions, and closing records without identifying and addressing the underlying causes of recurring operational problems.

This distinction matters. Investigation and problem solving are not the same activity. An investigation may establish what happened and meet procedural requirements. True problem solving seeks to understand why the problem occurred, what conditions made it possible, and what changes will prevent it from recurring.

Agentic AI creates an opportunity to improve not only the speed of investigations, but also the quality and consistency of manufacturing problem solving. Agents can assemble evidence, preserve operational context, surface relationships across systems, identify comparable events, and develop structured hypotheses. This allows subject-matter experts to spend less time gathering information and more time evaluating causes, testing assumptions, and defining effective countermeasures.

The examples below illustrate this broader manufacturing problem-solving pattern across production, quality, laboratory, engineering, and MS&T.

Deviation triage and investigation support

The agent prioritizes cases, builds an evidence pack, and links the event to batch, process, equipment, laboratory, and quality context. It can identify missing information, retrieve comparable events, propose root-cause hypotheses, and draft investigation content. Investigators determine which evidence and hypotheses are relevant and remain responsible for the final conclusion.

Batch release readiness and review by exception

The agent monitors whether required batch records, deviations, laboratory results, approvals, and supporting evidence are complete. It highlights unresolved exceptions and explains their potential effect on release readiness. Quality reviewers focus on high-risk or incomplete items rather than manually checking every record, while final release authority remains unchanged.

Operator and supervisor decision support

For frontline issues, the agent can interpret alarms and process conditions, retrieve approved procedures, show similar events, and recommend the next permitted step. Supervisors receive a shared view of production impact, open actions and escalation needs.

The design constraint at the frontline is severity of attention, not sophistication. An operator in gloves and goggles mid-batch will act on one clear permitted next step. They will not read a ranked list of five.

Supervisors need the opposite: breadth rather than focus. A shared view of production impact, open actions and what is about to require a decision.

Examples from manufacturing programs

AI-assisted deviation drafting has reduced manual preparation and shortened closure time for minor deviations. Predictive-maintenance solutions combine equipment, process and work-order information to identify degradation and explain recommended interventions. For example, in one program, condition monitoring built on FMEA-driven parameter selection and retrofitted vibration and temperature sensing addressed unplanned downtime and slow root-cause analysis. MES copilots help subject-matter experts retrieve recipe, batch and exception knowledge.

Case Study

Standardizing setup of critical equipment

Prescriptive analytics in the flow of work

Advanced analytics produced a prescriptive model from process data to standardize machine setup on critical equipment, where setup variation was propagating into product quality.

Baseline: 30–40 minutes and five to six setup repetitions per batch. Model output was expressed as recommended parameter values for the operator to apply, rather than as an analysis to be interpreted.

80% reduction in product failures arising from variation in process parameters.

Why it matters

Recurring losses that look small per batch are large per year.

Scaling intelligent operations

Where AI in this sector has produced measurable improvement, only a small minority of organizations have produced it at scale[2]. That is an industrialization gap, not a technology gap.

From isolated agents to agentic systems

A single agent supports a focused task; an agentic system coordinates specialist capabilities across an end-to-end process. Two design rules carry most of the weight: separate deterministic controls from probabilistic reasoning and make every hand-off between an agent and a person explicit. Ambiguous hand-offs are where accountability erodes.

Industrializing across sites

Enterprise scale depends on a global core with controlled local variation. Common semantics, integration adapters, risk controls, evaluation criteria, and user-experience patterns should be reused across sites. Deployment should be gated on measurable value, data readiness, validation effort and adoption. A site that fails the gate is not a site to be persuaded; it is a site whose prerequisites are not yet met.

Local teams configure site procedures, equipment, roles and regulatory requirements without rebuilding the solution. The test of scale is not how many agents exist but how quickly the next one can be built from components that already do.

Many organizations focus on the challenge of building the first successful agent. In practice, the more important question is how quickly the second, third, and tenth agents can be developed once the first is deployed. A single successful use case demonstrates technical feasibility. A repeatable development pattern demonstrates organizational capability.

This distinction is becoming increasingly important as manufacturers move beyond experimentation and begin to think about agent portfolios rather than individual agents. The organizations that realize the greatest value will not necessarily be those that create the most sophisticated initial solution. They will be the ones that establish reusable methods, common data foundations, clear governance models, and development approaches that allow new use cases to be delivered faster and with lower risk each time.

The real measure of industrialization is not whether a use case works. It is whether the next use case becomes easier to build because of what was learned from the first.

Enterprise Agentic platforms

Enterprise platforms provide the lifecycle capabilities needed to design, build, integrate and operate agents: an agent catalogue, workflow orchestration, connections to data products, identity and policy controls, evaluation, observability, and change management. Their purpose is specifically to avoid generating a new estate of point solutions.

The intelligent pharma operation

Not an autonomous factory without people, but a decision-centric operation in which trusted data, digital workflows, domain expertise, and AI work together. The defining characteristic is where quality sits: embedded in execution, evaluated continuously as the batch proceeds, rather than applied as a final checkpoint after the opportunity to intervene has passed.

This requires roles that do not exist in most organizations today: owners for individual agents, owners for data products, a model-risk function for manufacturing, and someone accountable for the design and performance of human–agent interaction.

People shift from searching for information and coordinating routine tasks to reviewing evidence, managing exceptions, and improving processes. Quality becomes a strategic decision enabler rather than a document reviewer.

Platform

A repeatable path to production

Capgemini's RAISE approach illustrates this platform model through four connected capabilities: accelerating reuse, building enterprise-grade agents, integrating multi-agent workflows, and operating agents in production.

The emphasis is on interoperability, continuous evaluation, policy enforcement, monitoring, and transparent human–agent interaction.

Whether through RAISE or an equivalent foundation, the requirement is the same: a repeatable path from proof of concept to controlled production. Organizations that lack such a path do not fail to build agents. They fail to keep them.

Enabling the digital manufacturing roadmap

Two failure modes bracket this work. One is waiting for a complete enterprise architecture before delivering anything. The other is launching disconnected pilots with no path to production. Each stage should deliver measurable value and strengthen the reusable foundation.

Horizon	Objective	Exit criteria before proceeding
0–3 months	Align on the first governed decision	Named business owner, agreed baseline; documented context of use, control envelope and risk classification; define decision rights
3–9 months	Prepare the semantic foundation and deliver one lighthouse use case in the flow of work	Measurable improvement against baseline; reusable data products identified; validation approach proven
9–18 months	Industrialize: governance, monitoring, support model, reusable agent components	Second use case delivered materially faster than the first, using the same core
18 months+	Scale across sites via a global core with controlled localization	Per-site deployment gated on value, data readiness, validation effort and adoption

Way Forward

Where to start: three conditions

Begin with a high-value decision or workflow where evidence is fragmented, response time matters, and human review can be clearly defined. Those three conditions together are what make a use case both valuable and governable.

Recommendations for manufacturing leaders

- Treat MES, data, and AI as one operating transformation, not three technology programs.
- Fund the decision, not the model.
- Build validation, cybersecurity and human accountability into the initial design[3].
- Standardize semantics before multiplying use cases[1].
- Establish joint ownership across Manufacturing, Quality, IT, Data and Engineering, and fund adoption beyond go-live.
- Scale only when value, trust and control effectiveness are demonstrable.

Conclusion: MES, from execution to intelligence

MES provides the trusted operational context and governed execution that make manufacturing AI relevant, reliable, and actionable. A Unified Namespace and data platforms extend that context across systems and sites, and agentic AI converts it into evidence, recommendations, and coordinated action.

None of this requires a new system of record or removing people from regulated decisions. It requires giving the intelligence something reliable to reason about and giving the organization a governed path from recommendation to action.

The opportunity is not to replace MES or human judgement. It is to build on existing investments and close the loop between manufacturing events, enterprise knowledge, governed decisions, and measurable outcomes. Organizations that combine strong execution foundations with reusable data, responsible governance, and scalable agent platforms can move beyond isolated pilots and establish intelligent operations as a durable enterprise capability.

The strategic shift is straightforward to state and demanding to execute: MES evolves from a system that records execution into the foundation that enables intelligence to act responsibly.

For more than a decade, pharmaceutical manufacturers have invested heavily in MES and related operational technologies. These investments have created trusted execution records, eliminated large amounts of manual activity, improved compliance, and generated vast quantities of operational data. In many respects, the industry has already built the foundation it needs.

The opportunity now is to capitalize on that investment. Advances in connectivity, contextualization, semantic architectures, and agentic workflows make it possible to use manufacturing information in ways that were previously impractical. The goal is no longer simply to record what happened, but rather to apply what has been learned to improve performance, solve problems more effectively, and make better decisions at the point where they matter.

The next era of manufacturing value realization will come not from generating more data, but from making better use of the data manufacturers already possess.

The industry has spent several years asking whether AI belongs in regulated manufacturing. As of January 2026, with ten jointly published principles on the table, that question has been answered[3]. The remaining question is one of sequencing.

The immediate leadership decision is not whether to pursue AI, it is where to establish the first governed information-and-action loop.

Authors

Sandeep Bhutani

Senior Director,Global Life Sciences Smart Factory, Capgemini
sandeep.bhutani@capgemini.com

Brian Eden

Vice President, Global Life Sciences Technical Operations, Capgemini
brian.eden@capgemini.com

Valerio Di Giovanni

Director, Manufacturing, Capgemini
Valerio.digiovanni@capgemini.com

Laurent Samot

Vice President, Digital Manufacturing Leader, Capgemini
laurent.samot@capgemini.com

To discuss where AI can deliver measurable value in your manufacturing and quality processes, contact any of the authors.

Acknowledgements:

This point of view draws on delivery experience across Capgemini's global life sciences manufacturing practice. The case studies are presented in anonymized form; client names and identifying details have been removed.

References

Figures quoted in this paper are attributed to their original source. Where a source is illustrative rather than measured, or where its scope or vintage limits how it should be read, that limitation is stated.

[1]Colley, James (BioPhorum). *"Data Issues and Regulatory Hurdles Are Stopping AI from Driving Real Value in Pharma."* BioProcess International 24(7), 2026: 240709. 22 July 2026.

[2]Deloitte US Center for Health Solutions and Deloitte Global. *"Confidence under pressure: How life sciences leaders are recalibrating for the rest of 2026."* 25 June 2026. Survey of 150 C-suite and senior life sciences executives, fieldwork April 2026.

[3]U.S. Food and Drug Administration - *FDA* and European Medicines Agency - *EMA*. Guiding Principles of Good AI Practice in Drug Development. January 2026.

[4]Dhage, Sarita Santosh. *"MES in Pharmaceutical Manufacturing: Enabling the Future of Personalized Medicine."* European Journal of Computer Science and Information Technology, 24 May 2025.

[5]Amery, Austin (EY). *"How electronic batch records can improve pharma manufacturing."* 10 April 2023. Figures cited are illustrative of review-by-exception mechanics rather than measured outcomes.

[6]McKinsey & Company. *"Digitization, automation, and online testing: Embracing smart quality control."* 14 April 2021. Estimates based on McKinsey models leveraging proprietary pharmaceutical manufacturing and quality benchmarks; scope is quality control operations.

[7]Boston Consulting Group. *"The AI-First Biopharma Company."* January 2026. Figures are BCG's statement of observed and achievable client impact; no sample size or methodology is published.

Further reading

- U.S. Food and Drug Administration, *"Artificial Intelligence in Drug Manufacturing"*, Discussion Paper, 2023.
- U.S. Food and Drug Administration, *"Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products"*, Draft Guidance, January 2025.
- European Medicines Agency, *"Reflection Paper on the Use of Artificial Intelligence in the Medicinal Product Lifecycle, 2024."*
- International Society of Automation, *"ANSI/ISA-95 and IEC 62264: Enterprise-Control System Integration."*
- National Institute of Standards and Technology, *"AI Risk Management Framework (AI RMF 1.0)"*, NIST AI 100-1, 2023.
- *"Expanding the Strategic Power of MES in Pharma"*, Capgemini, 2025.
- *"The Future of Pharma MES"*, Capgemini, 2024.
- *"MES of the Future: Insights from MES Berlin"*, Capgemini, 2024

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of over 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2025 global revenues of €22.5 billion.

Make it real.
www.capgemini.com