WHITE PAPER

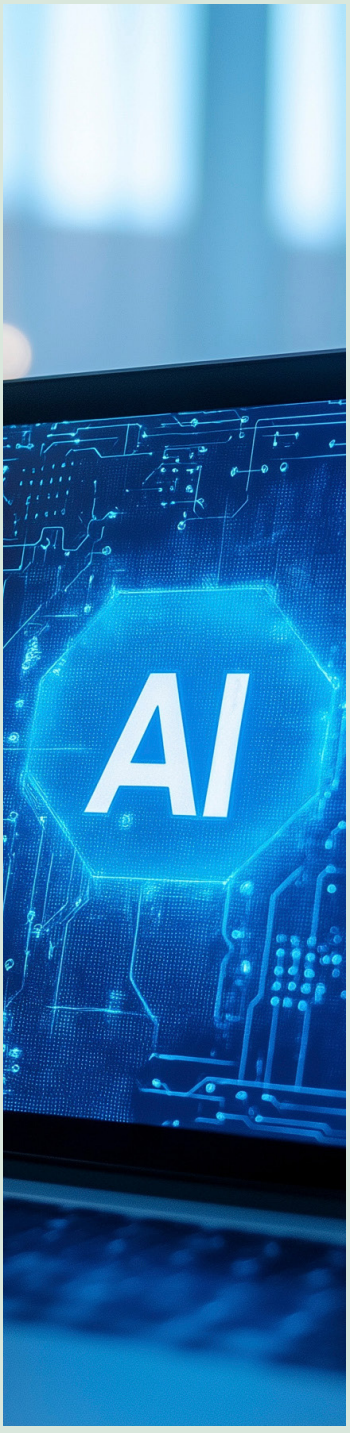

QUALITY BY DESIGN MEETS AI: REIMAGINING PLM-DRIVEN QUALITY ACROSS THE DIGITAL THREAD

AUTOMATING FMEA, NCR, RCA & CORRECTIVE / PREVENTIVE ACTION ACROSS THE DIGITAL THREAD

Infosys
Navigate your next

Reframing the product lifecycle as a closed-loop, intelligent quality system — where engineering, manufacturing, supplier and field data converge on a single digital thread, and where Quality by Design (QbD) supplies the scientific foundation that makes AI reasoning explainable, governed and continuously improving.

Executive Summary

Quality is no longer a downstream inspection gate — it is a lifecycle property that is designed, manufactured, and validated end-to-end. Yet the core processes that govern it — FMEA, NCR, RCA and CAPA — remain largely manual, document-centric, and siloed, living across spreadsheets, disconnected systems, and individual expertise. The result is a structural disconnect: engineering, manufacturing, supplier and field insights rarely loop back effectively into design decisions.

This paper argues that PLM is the natural backbone for AI-powered quality, because it already governs engineering definition, the digital thread, and the change process. By layering AI, knowledge graphs, and digital twins onto PLM — while federating MES, ERP, LIMS and field data — enterprises can shift quality from reactive compliance to a closed-loop, predictive, and ultimately autonomous capability.

Quality by Design (QbD) is introduced as the scientific foundation for this shift. QbD embeds quality into product and process design by explicitly defining critical quality attributes, critical process parameters, material characteristics, a design space, and a control strategy. When these QbD artifacts are governed in PLM and enriched by AI, quality knowledge becomes reusable, explainable, and continuously improving across the lifecycle. The regulated life-sciences sector — where QbD is codified in ICH Q8/Q9/Q10 — is the most mature instance of this pattern and proving ground for wider industrial application.

The four use cases reflect the lifecycle of quality. AI-driven FMEA leverages historical data to refine risk models and connect risks to QbD critical attributes. NCR automation enables real-time defect classification and design-space deviation detection. RCA applies graph-based reasoning for explainable insight across design, process, material, supplier and field evidence. Intelligent CAPA links actions to engineering changes, updates control strategies, and feeds lessons back into design — closing the quality loop.

The transformation in one line

From detecting defects after occurrence to predicting failures before release; from fragmented documentation to automated classification and cross-domain reasoning; and from isolated root-cause analysis to a continuously learning system integrated with the design of record — all under governed, auditable control.

1. Introduction

1.1 The evolution of quality management

Industrial quality management has evolved through distinct phases — from inspection (sorting defects post-production) to statistical process control, which shifted focus to processes. Later, Total Quality Management and Six Sigma embedded quality into systems and culture, supported by standards such as ISO 9001, AS9100, IATF 16949, and ISO 13485.

Today, quality is defined by the digital thread — the connected data flow linking requirements, design, manufacturing, and in-service performance. Quality is no longer a stage but a continuous, observable property across the lifecycle. While frameworks and standards exist, the persistent gap has been the lack of connected data and analytical capability — a gap AI is now positioned to close.

1.2 Quality by Design as the missing foundation

Quality by Design shifts the focus from detecting and correcting defects to designing robust products and processes from the beginning. Rather than treating quality as a downstream inspection or compliance activity, QbD defines the intended quality outcome, identifies what is critical to quality, understands the process variables that influence those outcomes, and establishes a control strategy that keeps product and process within an acceptable operating envelope.

In life sciences, QbD is expressed through constructs such as the Quality Target Product Profile (QTPP), Critical Quality Attributes (CQAs), Critical Process Parameters (CPPs), Critical Material Attributes (CMAs), design space, control strategy, and lifecycle process understanding — codified in ICH Q8(R2), Q9 and Q10. In broader industrial contexts the same logic maps naturally to critical-to-quality characteristics, key product characteristics, special characteristics, process controls, tolerance windows, parameter limits, control plans, and change-governed engineering definitions.

When QbD artifacts are managed inside PLM they become part of the product digital thread. This gives AI a high-quality reasoning foundation: not only what failed, but why it mattered, which characteristic was affected, which process parameter influenced it, which control was expected to prevent it, and which engineering change can permanently eliminate recurrence.

1.3 PLM as the system of record for engineering and quality

PLM platforms — such as Teamcenter, Windchill, 3DEXPERIENCE, and Aras Innovator — serve as the authoritative system of record for engineering definition: part master, bill of materials (BOM), CAD geometry, specifications, requirements, and controlled change processes (ECR/ECO). This uniquely positions PLM at the core of quality. The as-designed configuration, critical engineering characteristics, supplier of record, and change history all reside here; a Design FMEA traces back to requirements, while a CAPA-driven change is executed as an ECO within PLM.

Because PLM governs configuration and change, it anchors the lineage between as-designed, as-built, and as-maintained states. MES manages execution, ERP handles transactions, LIMS captures test and analytical data, and service systems capture field data — but PLM owns the underlying definition they all reference. That makes it the ideal backbone for closing the quality loop.

1.4 The limits of traditional, manual quality processes

Despite mature PLM and QMS investments, the day-to-day execution of FMEA, NCR, RCA and CAPA remains stubbornly manual. FMEAs are authored in spreadsheets and rarely updated after launch. CAPAs are tracked in a quality module disconnected from the engineering change that fixes the problem. NCRs are free-text records whose categorization depends on the individual writing them. Root-cause analysis relies on the tacit knowledge of a handful of experts and a whiteboard. The consequences are predictable: slow resolution, repeated failures, weak traceability, and quality knowledge that evaporates when people change roles. The remainder of this paper sets out how an AI layer on PLM dissolves these limits.

2. Industry Terminology & Foundations (PLM-Centric View)

This section defines the key quality constructs — QbD, FMEA, CAPA, NCR, RCA, and reliability frameworks — and their role in the digital thread. FMEA anticipates failure modes and risks; CAPA ensures issues are corrected and prevented through controlled actions; NCRs capture real-time quality deviations; and RCA identifies their underlying causes using structured methods. Reliability frameworks extend quality into in-service behavior, feeding field insight back into design and manufacturing. Together these constructs form a closed-loop quality system — but today remain weakly connected, limiting full traceability and reuse.

2.1 Digital thread and digital twin in a quality context

The **digital thread** is the connected, traceable flow of data linking requirements → design → manufacturing → service. In a quality context it lets a field failure on a specific serial number be traced back through its as-built record, its as-designed configuration, its FMEA, and the supplier of its failing component. The **digital twin** is a live, model-based representation of a specific asset or process, continuously updated with operational data. A process twin can simulate how a parameter drift propagates into a defect; a product twin can predict remaining life and surface emerging failure modes. Together they convert the digital thread from a static genealogy into a dynamic, predictive quality instrument.

Construct	What it answers	Where it lives	Primary digital-thread link
QbD / Critical Quality Knowledge	What quality outcome is intended, what is critical, and what controls keep the design/process robust?	PLM, QMS, process development, regulatory/technical files	Requirement → critical characteristic/CQA → CPP/CMA → design space/control strategy → control plan
DFMEA / PFMEA	How could this design or process fail, and how risky is it?	PLM (linked to part / process)	Function → requirement → control-plan characteristic
CAPA	How do we fix this and stop it recurring?	QMS / PLM quality module	Root cause → ECR/ECO → design or process change
NCR	What specifically failed, on which unit, and what is the disposition?	MES / QMS	Lot/serial → as-built → part → supplier
RCA	What is the true underlying cause?	Investigation (often ad-hoc)	Cross-domain: design × process × supplier × field
Reliability (SMRP/ RCM)	How does it behave and degrade in service?	EAM / FRACAS / field systems	As-maintained→ as-built → as-designed

The core quality constructs mapped to their system of record and digital-thread linkage.

2.2 The closed-loop quality instrument set

In a real PLM/QMS environment the abstract loop is realized as a defined set of controlled item types, governed end to end by a single **Corrective Action Plan (CAP)** that links them into one traceable record. The proactive side — governed by APQP / PPAP discipline — produces the FMEAs, control plans, and process-flow diagrams before launch; the reactive side acts when something escapes that plan. Engineering change is governed by CMII principles as a Problem Report → ECR → ECN/ECO chain, and process-origin fixes route through manufacturing process planning (MPP) so that work instructions, EBOM, and MBOM stay synchronized.

Closed-loop stage	Standard PLM/QMS item types	Purpose
Quality design intent	QTPP / CTQ / CQA, CPP, CMA, design space, control strategy	Defines what quality means and how it is achieved by design.
Quality event	Problem Report (PR), NCR, Audit Finding	Records that a problem exists — PR for field/customer, NCR for shop-floor non-conformance, Audit Finding for compliance gaps.
Containment	Hold, Stop-Ship, Purge Notice	Stop the bleeding while the investigation runs.
Analysis	RCA (5 Whys, Fishbone, FTA) + Cause	Establish the true underlying cause.
Corrective action	Deviation, Rework Order, Waiver	Disposition parts already affected — temporary, current production.
Preventive action	DFMEA/PFMEA update, control plan, Lessons Learned, ECO	Durable change that stops recurrence — future production.
Umbrella + change	Corrective Action Plan (CAP); CMII (PR → ECR → ECN/ECO); MPP	Links all of the above and drives the engineering / process change.

Closed-loop stages mapped to the standard PLM/QMS item types that realize them, all linked under one CAP.

2.3 QbD-to-quality use case mapping

Use case	How QbD adds value	AI enablement
AI-driven FMEA	Links failure modes to critical attributes, process parameters, material attributes and design-space assumptions.	Predicts risk priority, suggests controls, identifies missing failure modes, and updates living risk models.
NCR automation	Classifies defects against critical characteristics and detects deviations from approved parameter/process windows.	Auto-classifies NCRs, identifies emerging trends, and prioritizes high-risk events.
RCA acceleration	Provides a structured cause network across design intent, process parameters, material attributes and supplier evidence.	Uses knowledge graphs and explainable AI to propose evidence-backed causal hypotheses.
Intelligent CAPA	Ensures preventive actions update design space, control strategy, FMEA, control plan and engineering/process definitions.	Recommends controlled changes, validates effectiveness, and feeds lessons learned into future products.

One distinction matters throughout this paper. **Corrective** actions (Deviation, Rework Order, Waiver) are deliberately temporary — they disposition the parts already in production and are bound to the current lot or model year. **Preventive** actions (the FMEA and control-plan update plus the engineering change that implements it) are the durable fix for future production. A well-run closed loop keeps both, linked under the same CAP, so short-term disposition and long-term change never drift apart.

Value flow

Change occurs (design / process / supplier) → AI understands context → AI predicts risks (semantic expansion) → AI quantifies impact (cost, yield, quality) → system takes governed action.

3. Current Challenges in Quality Management

The limitations described above are not isolated annoyances; they compound into a systemic inability to learn. Seven interlocking challenges are observed consistently across automotive, aerospace, semiconductor, life-sciences and industrial-equipment enterprises.

3.1 Siloed systems: the PLM–MES–ERP–LIMS disconnect

PLM holds the as-designed definition, MES holds the as-built execution, ERP holds the transactional and supplier view, and LIMS holds analytical and test results. These systems were implemented at different times by different teams with different data models. A failure mode in the DFMEA (PLM) has no machine link to the defect record on the line (MES), the supplier scorecard (ERP), or the out-of-specification result (LIMS). Engineers, quality, and operations each see one face of the elephant, and reconciliation is manual.

3.2 Manual, reactive workflows

FMEAs are completed to satisfy a launch gate and then frozen. NCRs are dispositioned one at a time. RCA convenes only after a problem is large enough to demand attention. The dominant mode is reactive — the organization responds to escapes rather than anticipating them — and the labor cost of the manual workflow itself crowds out time for prevention.

3.3 Delayed feedback loops from shop floor to engineering

When the line discovers that a tolerance is unmanufacturable or a feature is failing, that signal must travel back to engineering to change the design. In practice this loop is measured in weeks and mediated by meetings and emails rather than structured data. By the time engineering acts, thousands of non-conforming parts may already exist.

3.4 Poor traceability across the lifecycle

Reconciling as-designed vs as-built vs as-maintained is the central traceability problem of modern manufacturing. Without it, a field failure cannot be reliably tied to the specific design revision and manufacturing conditions that produced it, recalls are scoped too broadly (or too narrowly), and the effectiveness of a corrective action cannot be proven.

3.5 Limited reuse of historical quality knowledge

Every FMEA, CAPA, and 8D report is a captured lesson — yet these lessons sit as unstructured documents that no one queries. The same failure mode is rediscovered on the next program; the same supplier defect recurs at the next plant. Knowledge that should compound instead decays.

3.6 Weak QbD-artifact governance

Many organizations define critical characteristics, control plans and process windows, but these artifacts are not always governed as living lifecycle knowledge. When QbD knowledge is disconnected from PLM change control, a CAPA can fix the immediate symptom without updating the underlying design intent, control strategy, or risk model — so the same failure re-emerges under a slightly different guise.

3.7 Data quality and standardization issues

Underlying all the above is inconsistent data: free-text defect descriptions with no controlled vocabulary, part numbers that differ between systems, missing or duplicated records, and FMEAs whose terminology varies by author. AI cannot reason reliably over data this noisy, which is why data governance is a precondition — not an afterthought — of the transformation.

Net effect

These challenges combine into an organization that cannot learn at the speed it operates. Quality knowledge is generated continuously but captured poorly, scattered across silos, and rarely fed back to the point of decision. The cost surfaces as warranty claims, scrap and rework, slow change cycles, and audit findings — the very metrics AI-powered quality is designed to move.

4. Why AI Is Needed in PLM-Driven Quality

AI is not a solution looking for a problem here. Four structural forces make intelligent automation a necessity rather than a luxury for industrial quality.

4.1 Increasing product complexity

Products are now cyber-physical systems combining mechanical, electrical, and software content, with configurations that multiply combinatorially. A modern vehicle, aircraft system, or semiconductor process has far more potential failure-mode interactions than any FMEA team can enumerate manually. Software-defined products also create cross-domain failures — a mechanical symptom with a firmware root cause. The analytical surface has outgrown human-only methods.

4.2 The need for predictive and proactive quality

Detecting defects after they occur is intrinsically expensive: the cost of a quality escape rises by roughly an order of magnitude at each stage it advances — from design to manufacturing to field. The economic imperative is to move detection upstream and to predict failure before it manifests. Prediction is a machine-learning problem; the data to support it (historical failures, process parameters, field returns) already exists but is unexploited.

4.3 The demand for a closed-loop digital thread

The digital thread generates enormous volumes of data: design parameters, manufacturing process traces, test results, sensor telemetry, warranty claims, and supplier records — far beyond what manual review can absorb. AI is the only practical means of finding the weak signals (the subtle correlation between a process parameter and a field failure) hidden in data at this scale.

Force	Why human-only methods fail	What AI contributes
Product complexity	Too many failure-mode interactions to enumerate	Exhaustive pattern search across designs and history
Predictive quality	Detection happens after the fact	Failure prediction before release and before escape
Data scale	Volume exceeds manual review capacity	Weak-signal detection across the full thread
QbD lifecycle control	Critical attributes, process parameters and control strategies must stay aligned across changes	Automated impact analysis and recommended updates to FMEA, control plans and design-space assumptions
Closed-loop thread	Cross-domain reasoning is slow and manual	Near-real-time correlation and feedback to design

The structural forces driving AI adoption in quality, and the specific capability AI contributes to each.

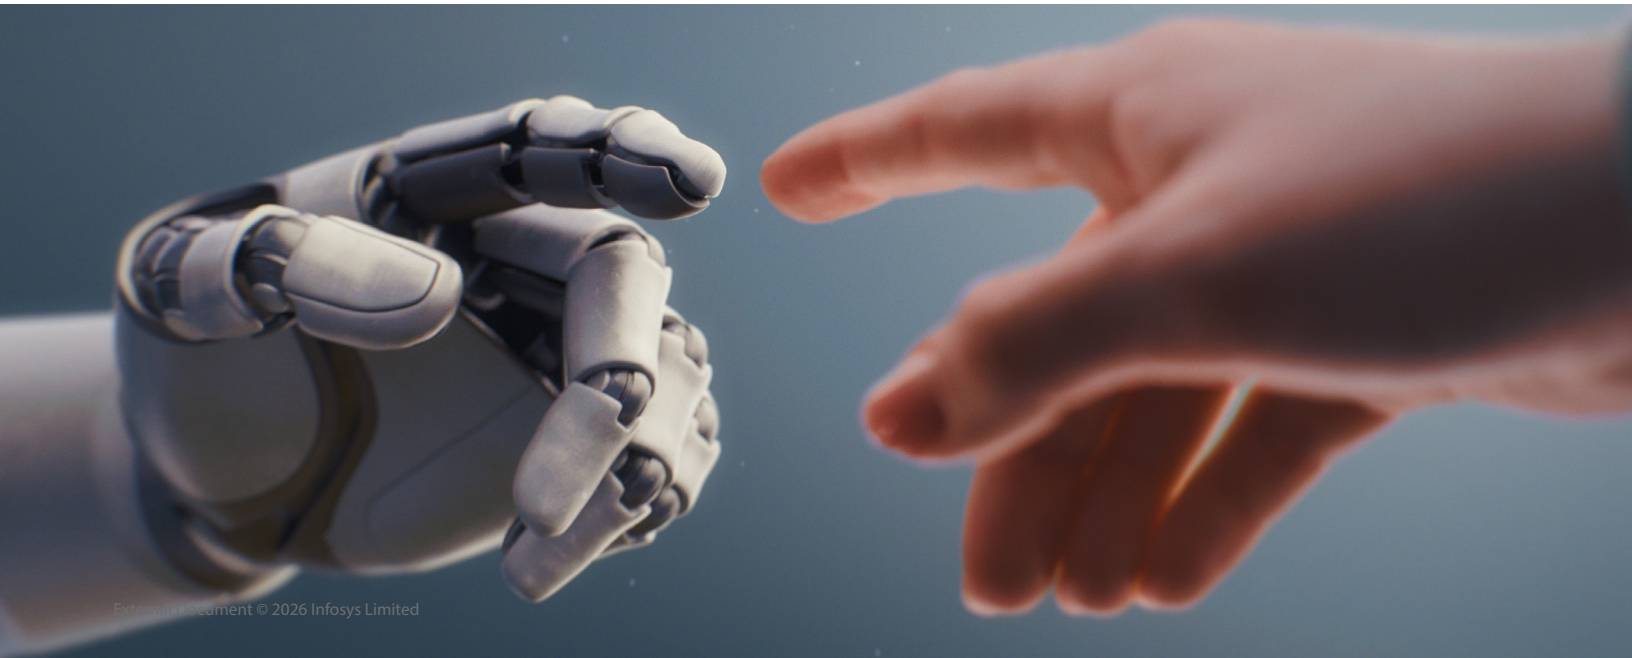

5. AI-Powered Quality: Conceptual Architecture

This section describes a reference architecture for AI-powered quality, built on a layered model with PLM as the backbone. The design principle throughout is federation, not replacement: existing systems remain masters of their own data; the architecture adds a connective and intelligent layer above them.

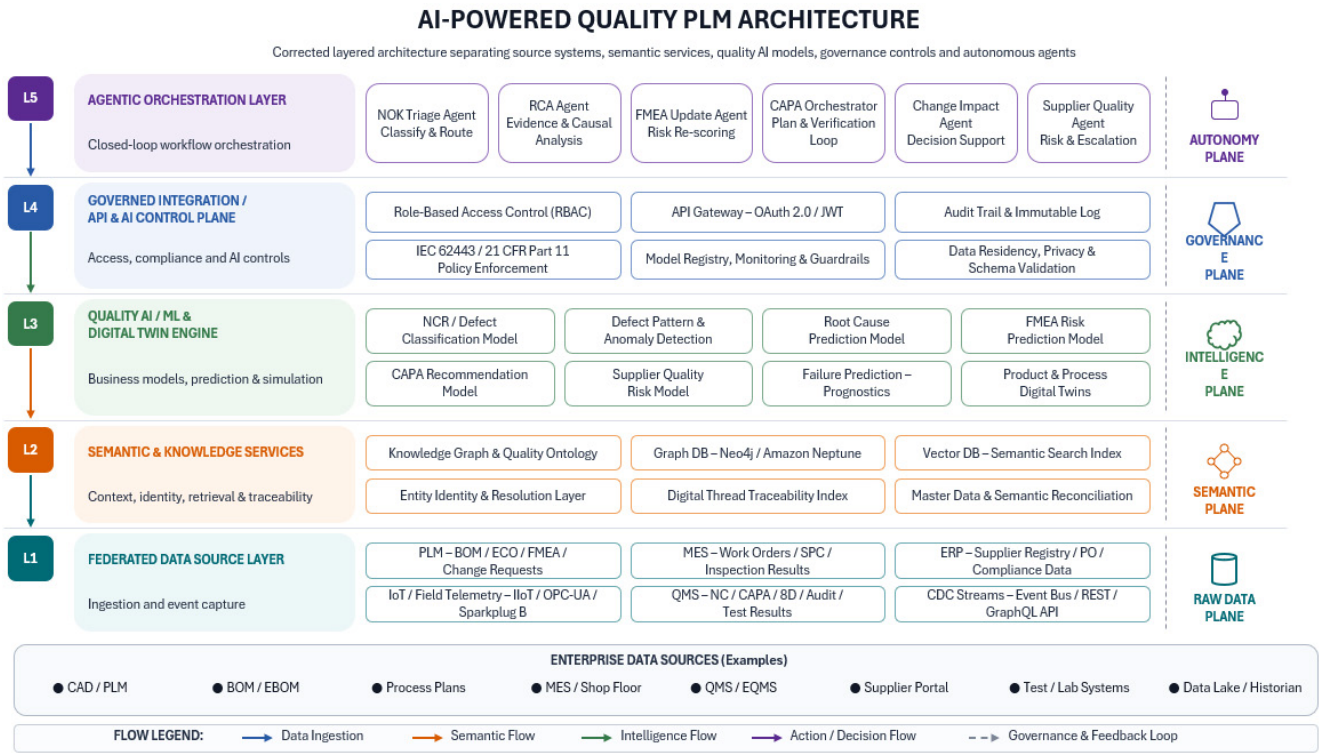

AI-Powered Quality PLM reference architecture — a layered intelligence stack from federated data sources through the semantic, intelligence, governance and autonomy planes.

5.1 The layered reference model

The architecture is best understood as cooperating layers, loosely aligned to the ISA-95 automation hierarchy so that data context is preserved from the shop floor to the enterprise:

- **QbD foundation layer.** Critical quality knowledge, critical characteristics, CPP/CMA relationships, design space, control strategy, risk assumptions, and lifecycle process understanding.
- **System-of-record layer (PLM).** The authoritative engineering definition: parts, BOMs, requirements, specifications, FMEAs, control plans, and the ECR/ECO change process — the anchor to which all quality intelligence is tied.
- **Operational data layer (MES, ERP, LIMS, IoT/field).** As-built genealogy, process parameters, NCRs, supplier and transaction data, analytical results, and in-service telemetry and reliability data.
- **Integration and semantic layer.** A federated data fabric that normalizes identifiers (part, lot, serial, supplier), harmonizes vocabularies, and maps records to a common ontology — where as-designed, as-built, and as-maintained are reconciled.
- **AI/ML and Business Model layer.** The analytical core: pattern detection, failure prediction, classification, knowledge mining, and the knowledge graph encoding relationships among failures, causes, parts, processes, and suppliers.
- **Decision and experience layer.** Where intelligence reaches people and systems: AI-assisted FMEA, NCR, RCA and CAPA workflows surfaced inside the PLM and quality tools engineers already use, plus dashboards and — increasingly — autonomous agents.

5.2 PLM as the backbone

PLM is the backbone for three reasons. First, it owns configuration and change — the only credible anchor for lifecycle traceability. Second, it owns the engineering semantics (functions, requirements, characteristics) that give quality data its meaning. Third, it owns the change workflow through which any corrective action must ultimately be implemented. Placing the intelligence around PLM ensures that insight lands where decisions are made and changes are controlled, rather than in a disconnected analytics sandbox.

5.3 The three-layer QbD × PLM × AI integration stack

Beneath the reference model sits a tighter conceptual stack that explains why the whole system is trustworthy. Each layer depends on and enriches the one below it, with bidirectional data flows so that every AI prediction is grounded in PLM-managed process data, and every PLM workflow is governed by the QbD design space.

Layer	Role and constituent capabilities
AI intelligence layer	Predictive analytics, NLP-driven document assembly, ML-based process control, autonomous CAPA triage, and knowledge graphs. Reasons over the layers below and proposes governed action.
PLM integration framework	Change management, document control, BOM linkage, stage-gate workflows, regulatory filing, and version control. Turns insight into controlled change.
QbD foundation	QTPP, CQA definition, CPP/CMA, design space, control strategy, and process understanding (ICH Q8/Q9/Q10 in life sciences; CTQ/KPC/control plans in discrete industry). Supplies the meaning AI reasons over.

The QbD × PLM × AI stack. Grounding every AI prediction in PLM-governed QbD knowledge is what makes the intelligence explainable and auditable.

5.4 The AI/ML layer: three core capabilities

- **Pattern detection.** Unsupervised and supervised models cluster NCRs, surface recurring defect signatures, and flag anomalous process behavior — finding structure in high-volume, noisy quality data.
- **Failure prediction.** Models trained on historical failures, process parameters, and field returns estimate the probability and timing of failure modes, enabling proactive FMEA enrichment and predictive-maintenance triggers.
- **Knowledge mining.** NLP and large-language-model techniques extract structured knowledge from unstructured FMEAs, 8D reports, CAPAs, and service notes, making decades of tacit quality knowledge queryable and reusable.

5.5 Knowledge graphs and digital twins

A **knowledge graph** represents quality entities — parts, functions, failure modes, causes, processes, suppliers, NCRs, CAPAs, field events — as nodes, and their relationships as typed edges. This is the structure that makes cross-domain RCA possible: from a field-failure node a traversal can reach the as-built process, the responsible supplier, the design characteristic, and any prior similar failures, ranking causal hypotheses by the weight of connected evidence — and providing the explainable path that quality auditors require. **Digital twins** add a model-based, predictive dimension: a process twin simulates how a manufacturing-parameter drift propagates into a dimensional defect (validating the FMEA detection rating against simulation rather than estimate), while a product twin fed by field telemetry predicts remaining useful life and surfaces emerging failure modes that feed the next design iteration.

5.6 Agentic AI: from assistant to actor

The emerging frontier is agentic AI — goal-directed agents that not only surface insight but orchestrate multi-step quality workflows. An NCR-triage agent can classify an incoming defect, query the knowledge graph for similar history, draft a containment recommendation, and open a candidate CAPA — escalating to a human for the decisions that require judgment. Architecturally, agents operate at the decision layer and act through governed interfaces (a common pattern is to expose PLM, MES, and knowledge-graph capabilities to agents through a standardized protocol layer), always within the controls of the change process. Agentic AI is optional today and powerful tomorrow; the architecture should be agent-ready even where agents are not yet deployed.

Architectural principle — federate, don't replace

Each source system remains master of its data. The semantic layer reconciles identity and meaning; the knowledge graph encodes relationships; the AI layer reasons; and PLM remains the backbone where insight becomes controlled change. This keeps the transformation incremental, auditable, and resilient to future tool changes.

6. The AI Capability Catalogue

AI does not simply automate existing quality workflows — it transforms them. Nine distinct capabilities, drawn from mature deployments in regulated and discrete manufacturing alike, create a system that is predictive rather than reactive, continuously learning rather than static, and capable of synthesizing knowledge across products, sites, and time horizons that exceed human cognitive capacity. Each maps onto the reference architecture and is grounded in PLM-governed QbD knowledge.

Capability	What it does	Grounded in
Predictive design-space mapping	ML models trained on DoE / test data extrapolate the design space beyond experimentally tested points, predicting responses at unsampled parameter combinations with quantified uncertainty; Bayesian optimization finds the ideal operating point with fewer experiments.	QbD design space, DoE records
NLP for documentation & change	LLMs read standards, precedent filings, and current PLM state to draft investigations, 8D reports, and change-impact narratives, and to flag gaps — cutting authoring time while improving consistency.	PLM change records, FMEAs
Real-time process monitoring & control	Digital twins fed by process/PAT sensors predict quality attributes in near-real-time; when a parameter drifts toward the boundary, AI issues a corrective set-point recommendation before an out-of-spec event — enabling soft intervention instead of batch failure.	Control strategy, CPP ranges
Autonomous CAPA triage & linkage	Incoming deviations are auto-classified by severity and linked to their most probable root-cause node in the knowledge graph; AI recommends CAPAs ranked by historical effectiveness and checks for recurrence across the site network.	Knowledge graph, CAPA history
Risk-based change-impact scoring	When a change enters PLM, AI evaluates its impact across the design-space model, control strategy, and downstream documents simultaneously, generating a quantitative risk score and recommending the minimum required approval pathway.	Design space, control plan
Knowledge graph for lifecycle continuity	A graph links every development decision — rationale, failure modes, material-variability studies — to the PLM product record, so a similar question years later surfaces the relevant historical evidence, protecting institutional memory.	PLM product record
Formulation / process optimization	Generative models propose novel material combinations or process sequences optimized against target profiles; federated learning across products and sites identifies cross-product patterns no single team would discover alone.	QTPP / target profile
Supplier & material risk prediction	AI analyzes incoming material-variability data against QbD-established ranges and flags suppliers whose trends are converging on the boundary — enabling proactive qualification before disruption.	CMA ranges, supplier data
Regulatory / standards intelligence	AI monitors agency communications, findings, and inspection trends to identify signals relevant to a product's design space or control strategy — allowing the organization to self-correct before an external finding.	Standards corpus, audit data

Nine AI capabilities that amplify the QbD × PLM integration, each anchored to a governed source of truth.

7. Lifecycle Integration: From Concept to Retirement

The integration operates continuously across the entire product lifecycle, with each stage contributing to and drawing from a shared quality knowledge base. The table below traces the pattern using life-sciences terminology (where it is most formalized); the same stage logic applies to discrete-manufacturing programs, substituting APQP/PPAP gates and control plans for the regulatory constructs.

Lifecycle stage	What happens	AI contribution
Early development (QTPP / target-setting)	PLM captures the initial target profile as a structured record; the knowledge baseline against which all future decisions are logged is established.	Mines databases and literature to suggest critical-attribute candidates and preliminary ranges.
Formulation & process development	DoE experiments are PLM-linked to the parameter–attribute model; the emerging design space is stored as a multidimensional object, not a static table.	Bayesian optimization proposes the next run, reducing experiments by 30–50%.
Clinical / pre-production	The control strategy is formalized in PLM, linking each critical attribute to its detection method, parameter range, and in-process test.	AI-assisted risk assessment scores FMEA severity, occurrence, and detectability using process data and precedent.
Regulatory filing / launch gate	Authoring draws structured records from PLM; cross-references between development, process description, and controls sections are maintained automatically.	NLP drafts filing sections and compares the draft against guidances, highlighting gaps before submission.
Commercial manufacturing	Continual-verification dashboards are fed from PLM control-strategy limits; digital twins run in parallel with each batch/lot.	Deviations from the design space trigger AI-generated corrective actions before product is affected.
Lifecycle management (post-launch change)	A proposed change — supplier switch, equipment upgrade, process improvement — invokes the design-space model through the PLM change workflow.	If the change stays within the approved space, AI classifies it as a lower-tier regulatory item, saving months of filing time.
End of life / line extension	For a line extension or new variant, the full knowledge graph from the original product is retrieved and re-usable elements identified.	AI identifies what transfers, so re-development scope shrinks and the new product starts from a knowledge-rich baseline.

The QbD × PLM × AI integration across the lifecycle — each stage both enriches and draws from the shared knowledge base.

8. Worked Lens: Regulated Life Sciences

Life sciences is where QbD is not optional but codified — in ICH Q8(R2) (pharmaceutical development), Q9 (quality risk management), and Q10 (pharmaceutical quality system), reinforced by FDA process-validation and PAT guidance and EMA/CHMP QbD guidelines. That regulatory rigor makes it the clearest proving ground for the closed-loop pattern this paper describes. The central argument is stark: **QbD without PLM is science locked in documents; PLM without QbD is workflow without quality rationale; adding AI closes the loop** — converting the design space from a regulatory deliverable into a living operational asset that governs every decision across the product's commercial life.

8.1 Before and after — the quality dimensions transformed

Dimension	Without integration	With QbD × PLM × AI	Driver
Design-space management	Static spreadsheets; manual version control	Living multidimensional model in PLM; AI-predicted boundaries	AI + PLM
Change management	Document-based; no automatic risk scoring	Automated design-space impact check; risk-scored pathway recommendation	AI
Regulatory filing	Manual authoring; version mismatches common	PLM-driven assembly; AI gap analysis against guidances	PLM + NLP
Process monitoring	Retrospective SPC; batch-level review	Real-time digital twin; predictive attribute forecasting within design space	AI + PAT
CAPA management	Manual root-cause linking; paper trail	AI links deviation to CPP/CQA; recommends evidence-ranked CAPAs	AI
Knowledge continuity	Tacit knowledge lost at staff turnover	Structured knowledge graph persists in PLM; AI retrieval at reuse	PLM + KG

The regulated-quality operating model, before and after the integrated capability is in place.

8.2 Quantified performance improvements

Organizations implementing the integration report consistent improvements across several performance dimensions. The figures below are directional planning ranges reported by early adopters — not guarantees — and depend heavily on data readiness and integration depth.

Performance dimension	Reported improvement
Regulatory submission quality — first-cycle approval rate	+40%
Change-management cycle-time reduction	-60%
Batch-failure prediction accuracy (AI-assisted continual verification)	92%
CAPA cycle-time reduction with AI triage	-45%
DoE experiment reduction with AI-guided design	-35%
Re-development cost avoidance for line extensions	-38%

Directional impact ranges reported by early adopters of the QbD × PLM × AI operating model.

8.3 The change event, visualized

The diagram below traces a single quality-impacting change through the five-step closed loop this paper advocates — from a change event through AI contextualization, semantic expansion, federated impact quantification, and finally closed-loop execution — with a continuous learning loop feeding outcomes back to strengthen the models.

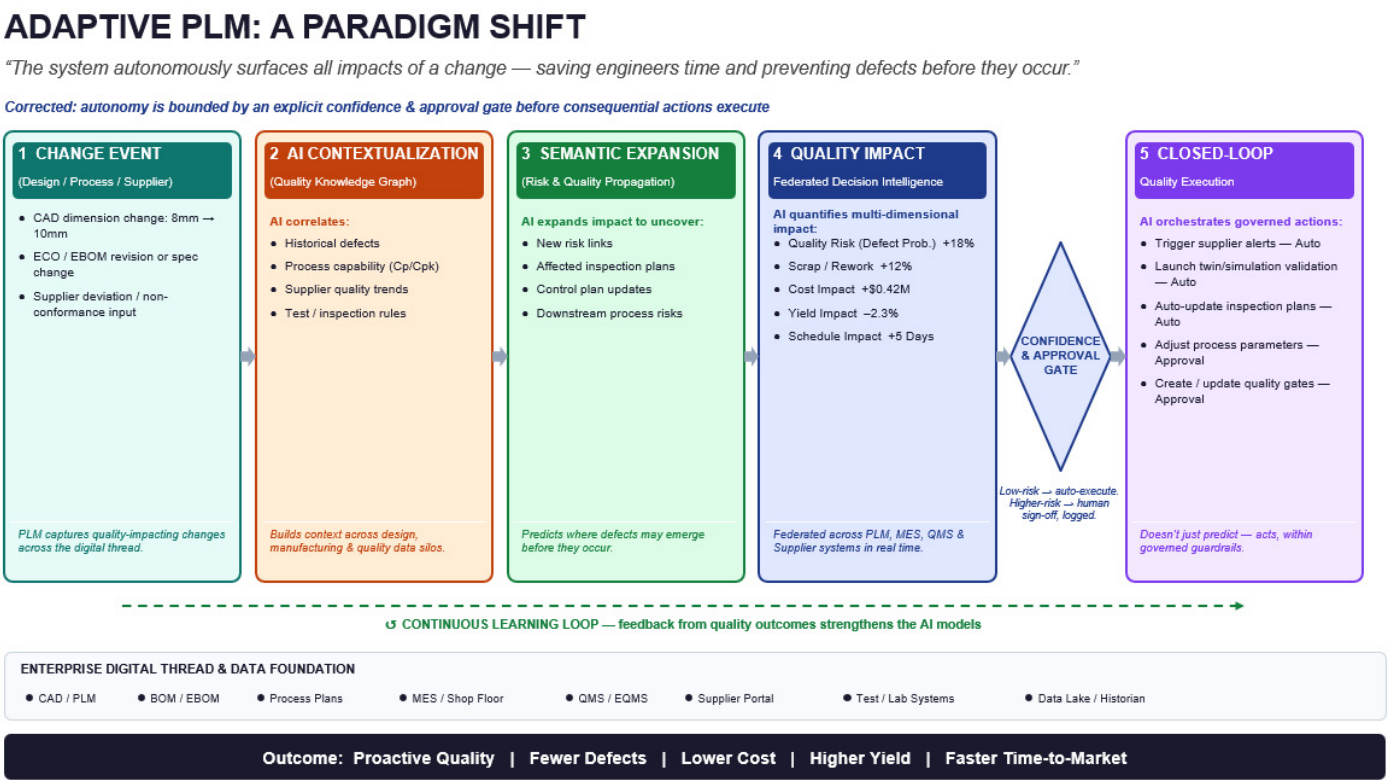

Adaptive PLM: a paradigm shift. The system autonomously surfaces all impacts of a change — saving engineers time and preventing defects before they occur.

9. Benefits and Business Impact

The value of AI-powered quality is concrete and measurable across operational, financial, and compliance dimensions. The ranges below reflect outcomes reported by early-adopter industrial enterprises; they are directional planning figures, not guarantees, and depend heavily on data readiness and scope.

Value driver	Mechanism	Typical reported impact
Quality resolution time	Auto-classified NCRs, AI root-cause suggestions, linked CAPAs	30–50% reduction in RCA/CAPA cycle time
First-time-right / yield	Predictive FMEA and early trend detection prevent escapes	10–20% improvement in first-pass yield
Engineering change cycles	Enforced CAPA–ECO linkage and faster feedback loops	20–40% faster change implementation
Knowledge reuse	Queryable knowledge graph of failures and resolutions	Repeated-failure recurrence materially reduced
Cost of poor quality (COPQ)	Fewer escapes, less scrap/rework, fewer warranty events	Significant multi-point COPQ reduction
Compliance & audit readiness	Complete, machine-traceable lineage across the thread	Audit-preparation effort sharply reduced
QbD lifecycle effectiveness	Design-space and control-plan updates linked to NCR/RCA/CAPA outcomes	Fewer avoidable deviations; stronger preventive posture

Value drivers and typical reported impact ranges. Actual results depend on data quality, integration depth, and process scope.

9.1 The strategic dividend

Beyond the line items, the strategic benefit is compounding organizational learning. Each resolved NCR, CAPA, and RCA enriches the knowledge graph, so the marginal cost of solving the next similar problem falls toward zero. Over time the enterprise builds a defensible quality asset — a living, queryable memory of everything it has learned about how its products fail and how to prevent it — that no spreadsheet archive can replicate and no departing expert can take with them.

10. Implementation Strategy: A Practical Roadmap

A successful program is incremental, data-first, and anchored on PLM. It does not begin with a model; it begins with the data and the loop that the model will eventually close.

10.1 Data readiness and governance

Data is the binding constraint. Before models, invest in a controlled defect taxonomy, consistent part/lot/serial/supplier identifiers across systems, and the semantic mapping that reconciles as-designed, as-built, and as-maintained. Establish data ownership, quality metrics, and stewardship. This is unglamorous and decisive: model performance is capped by data quality, and the governance built here is also what makes the AI auditable later.

10.2 Integration with legacy PLM/MES/ERP/ LIMS

Integrate through a federated fabric rather than a monolithic data-warehouse migration. Use the source systems' APIs and event streams, map to a common ontology, and materialize the knowledge graph from the federated layer. Align contextualization to ISA-95 so operational data carries its meaning upward. This preserves existing investments and avoids a multi-year, high-risk re-platforming.

10.3 AI model lifecycle management (MLOps)

- Versioning and reproducibility. Treat models, training data, and features as versioned, governed artifacts — essential in regulated industries where model decisions must be reproducible for audit.
- Monitoring and drift detection. Quality-data distributions shift as products and processes change; monitor for model drift and retrain on a governed cadence.
- Human-in-the-loop validation. Models propose; engineers decide. Capture engineer acceptance/rejection as feedback to improve the models and to maintain accountability.

10.4 Change management and adoption

The hardest barrier is rarely technical. Quality engineers must trust AI suggestions, which requires explainability, demonstrated accuracy on their own data, and positioning the AI as an assistant that removes drudgery rather than a replacement that removes judgment. Start with augmentation, show the time saved on a painful task, and let credibility compound. Executive sponsorship and a cross-functional team spanning quality, engineering, manufacturing, regulatory and IT/OT are prerequisites.

Phase	Focus	Representative outcome
Foundation	Taxonomy, identifiers, semantic layer, governance	Clean, federated quality data on the thread
First use case	NCR classification & trend detection	Early wins, training data, organizational trust
Knowledge graph	Connect design × process × supplier × field	RCA acceleration and reuse
Predictive	AI-FMEA and CAPA root-cause assistance	Proactive prevention; faster change cycles
Closed loop / agentic	Automated feedback to design; assisted agents	Self-improving, increasingly autonomous quality

Value drivers and typical reported impact ranges. Actual results depend on data quality, integration depth, and process scope.

11. Risks and Considerations

Deploying AI in quality — a domain where decisions carry safety, regulatory, and liability consequences — demands disciplined risk management. Four areas require explicit governance.

11.1 Data quality and bias

Models are only as good as their training data. Biased history — for example, defects systematically under-recorded on a particular line — produces biased predictions, and historical FMEA optimism can be learned and perpetuated. Mitigation requires the data governance of Section 10, active monitoring for bias, and human validation of high-stakes outputs.

11.2 Explainability of AI decisions

In quality, an unexplained recommendation is unusable: auditors, regulators, and engineers must understand why. This is the decisive advantage of graph-based reasoning and retrieval-grounded methods over opaque models — they yield an inspectable evidence path. Favor explainable techniques for decision-grade outputs, require every AI suggestion to carry its supporting evidence, and keep a human accountable for the decision. Explainability is not a nice-to-have here; it is a compliance requirement.

11.3 Cybersecurity (IEC 62443)

Connecting PLM, MES, and OT systems expands the attack surface across the IT/OT boundary. The architecture must be secured in line with IEC 62443, the standard for industrial automation and control-system security: zone-and-conduit segmentation between IT and OT, least-privilege access, secured integration interfaces and event streams, and protection of the quality data itself, which is sensitive intellectual property and regulatory evidence. Agentic components that can act on systems require especially tight authorization and audit controls.

11.4 Regulatory compliance

Regulated industries impose specific obligations on quality records and any system that influences them. 21 CFR Part 11 (life sciences) requires electronic-record integrity, audit trails, and validated systems; AS9100 (aerospace), IATF 16949 (automotive), and IEC 61508 (functional safety) impose traceability, control, and competence requirements. AI systems that touch CAPA, NCR, or FMEA records fall within scope and must be validated, version-controlled, and auditable. Engage quality and regulatory functions from the outset — retrofitting compliance onto a deployed AI system is far costlier than designing it in.

Governing principle for risk

AI proposes; humans decide; the system records why. Every AI output is explainable and evidence-linked, every decision retains a human owner, every action is audit-traceable, and the whole is secured to IEC 62443 and validated to the applicable regulatory standard. Under this principle, AI augments quality judgment without ever escaping its controls.

12. Outlook

The trajectory of AI-powered quality points toward systems that are progressively more autonomous — while remaining firmly within human and regulatory control.

12.1 Autonomous quality systems

As models mature and trust accrues, routine quality decisions — NCR disposition for well-understood defect classes, FMEA updates triggered by field data, low-risk preventive-action propagation — will be increasingly automated, with humans focusing on novel and high-severity cases. The role of the quality engineer shifts from executing workflows to governing an intelligent system and exercising judgment where it matters most.

12.2 Closed-loop, self-healing manufacturing

The ultimate expression of the closed-loop digital thread is self-healing manufacturing: a process digital twin detects an emerging drift, predicts the defect it will cause, and adjusts process parameters — or triggers a design or control-plan change — before any non-conforming part is produced. Quality becomes a continuous, real-time control loop rather than a periodic inspection, and the distinction between preventing and detecting defects largely dissolves.

12.3 Generative and agentic AI in PLM

Generative AI will draft FMEAs, 8D reports, and CAPA investigations from underlying data, compressing the documentation burden that consumes so much quality-engineering time. Agentic AI will orchestrate end-to-end quality workflows — triaging a defect, investigating its cause, drafting and linking the corrective action, propagating the preventive action across the fleet, and escalating to humans only for judgment calls. Standardized agent-to-system protocols will let these agents act safely across PLM, MES, and the knowledge graph within governed boundaries. The destination is a quality function that operates at machine speed and human scale simultaneously.

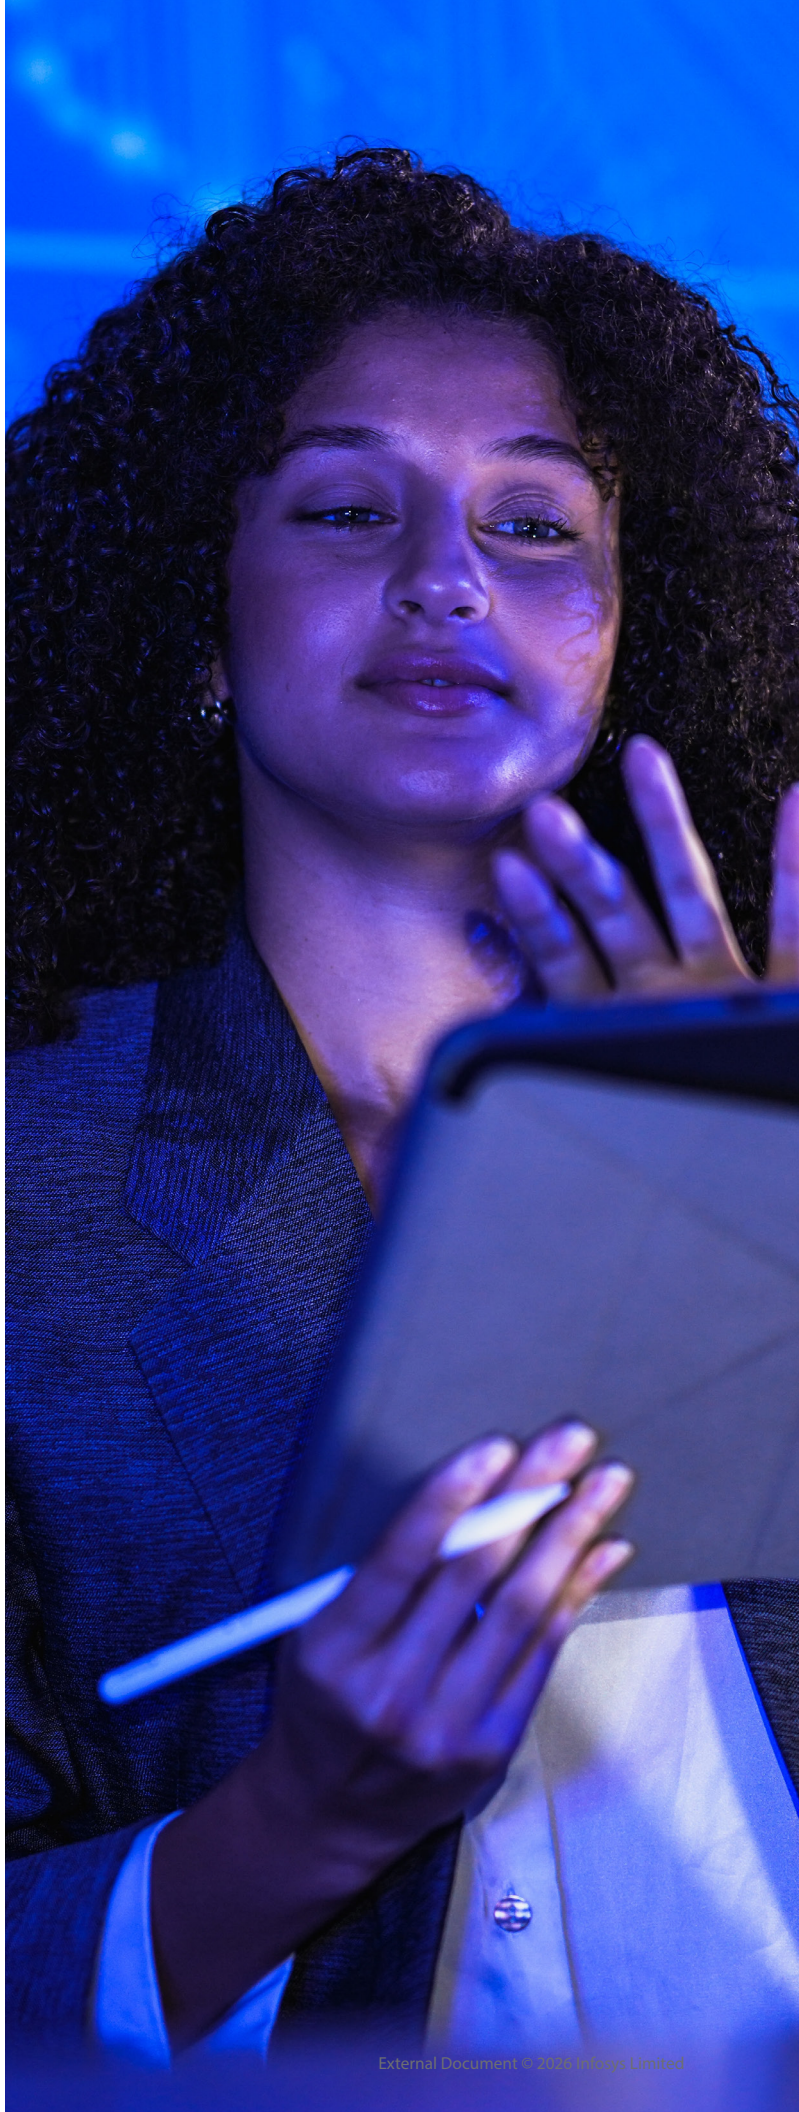

13. Conclusion

Industrial quality is at an inflection point. While frameworks like FMEA, NCR, RCA and CAPA remain robust, their manual, document-heavy execution cannot keep pace with rising product complexity and data scale — leading to escapes, slow change cycles, and decaying knowledge.

This paper argues that the solution largely exists. PLM provides the natural backbone for AI-powered quality by owning engineering definition, configuration, and change. By federating MES, ERP, LIMS and field data, modeling relationships in a knowledge graph, grounding reasoning in QbD, and applying explainable AI, enterprises can transform quality processes — turning FMEA into a living risk model, NCR into a real-time signal, RCA into reusable cross-domain intelligence, and CAPA into linked, verifiable change.

The technology is ready and the regulatory precedent is forming. The result is a shift from reactive to predictive — and ultimately autonomous — quality, where failures are anticipated, knowledge compounds, and the digital thread closes the loop from field back to design. **The strategic window to build this capability ahead of industry-wide adoption is now.**

About The Authors

Sibaji Pattanaik
Principal Consultant, IOT, Infosys

Kapil Manikrao Jadhav
Industry Principal, IOT, Infosys

14. References & Further Reading

The following standards, frameworks, foundational works, and technical sources inform the methodologies and architecture described in this paper. Standard designations reflect editions current at the time of writing; readers should confirm the latest revision before use.

Standards, Frameworks & Regulations

1. ANSI/ISA-95.00.01 (IEC 62264-1 Mod). Enterprise-Control System Integration — Part 1: Models and Terminology. ISA.
2. IEC 62264 (Parts 1–5). Enterprise-Control System Integration. IEC — the international equivalent of ISA-95, based on the Purdue Reference Model.
3. ISA/IEC 62443 Series. Security for Industrial Automation and Control Systems (incl. 62443-3-3 System Security Requirements and Security Levels; 62443-3-2 Security Risk Assessment for System Design). ISA / IEC.
4. AIAG & VDA FMEA Handbook (2019). Failure Mode and Effects Analysis: Design FMEA, Process FMEA, Supplemental FMEA for Monitoring & System Response. AIAG.
5. SAE J1739. Potential Failure Mode and Effects Analysis in Design (DFMEA) and in Manufacturing and Assembly Processes (PFMEA). SAE International.
6. ISO 9001:2015. Quality Management Systems — Requirements. ISO.
7. IATF 16949:2016. Quality Management System Requirements for Automotive Production and Relevant Service-Part Organizations. IATF.
8. SAE AS9100D (2016). Quality Management Systems — Requirements for Aviation, Space, and Defense Organizations. SAE / IAQG.
9. ISO 13485:2016. Medical Devices — Quality Management Systems — Requirements for Regulatory Purposes. ISO.

10. U.S. FDA, 21 CFR Part 11. Electronic Records; Electronic Signatures. U.S. Food and Drug Administration.

11. U.S. FDA, 21 CFR Part 820 / QMSR. Quality System Regulation / Quality Management System Regulation. U.S. FDA.

12. IEC 61508. Functional Safety of Electrical/Electronic/ Programmable Electronic Safety-related Systems. IEC.

13. ISO 31000:2018. Risk Management — Guidelines. ISO.

14. ICH Q8(R2), Q9, Q10. Pharmaceutical Development; Quality Risk Management; Pharmaceutical Quality System. International Council for Harmonisation.

15. FDA Process Validation Guidance (2011) & PAT Guidance (2004); EMA/CHMP QbD Guidelines; EU GMP Annex 11. Regulatory basis for QbD, continual process verification, and computerized-system compliance.

Quality Methodology & Foundations

16. Feigenbaum, A. V. (1991). Total Quality Control (3rd ed.). McGraw-Hill — company-wide TQC and cost-of-quality.

17. Deming, W. E. (1986). Out of the Crisis. MIT Press — the PDCA closed-loop improvement cycle.

18. Juran, J. M., & De Feo, J. A. (2016). Juran's Quality Handbook (7th ed.). McGraw-Hill.

19. Ishikawa, K. (1985). What Is Total Quality Control? The Japanese Way. Prentice-Hall — cause-and-effect (fishbone) analysis.

20. Ford Motor Company. Global 8D / Team Oriented Problem Solving (TOPS) structured problem-solving methodology.

21. Nowlan, F. S., & Heap, H. F. (1978). Reliability-Centered Maintenance. U.S. DoD — foundational RCM report.

22. Moubray, J. (1997). Reliability-Centered Maintenance (RCM II, 2nd ed.). Industrial Press.

23. SAE JA1011 / JA1012. Evaluation Criteria for, and a Guide to, Reliability-Centered Maintenance (RCM) Processes. SAE International.

24. SMRP. Society for Maintenance & Reliability Professionals — Body of Knowledge and Best Practices.

PLM, Digital Thread & Digital Twin

25. Grieves, M. (2014). Digital Twin: Manufacturing Excellence through Virtual Factory Replication. White Paper (with NASA's J. Vickers).

26. Stark, J. Product Lifecycle Management (Vols. 1–4). Springer.

27. Sääksvuori, A., & Immonen, A. Product Lifecycle Management. Springer.

28. NIST. Model-Based Enterprise (MBE) and Digital Thread for Smart Manufacturing program publications.

AI, Knowledge Representation & Governance

29. Hogan, A., et al. (2021). "Knowledge Graphs." ACM Computing Surveys, 54(4).

30. Lewis, P., et al. (2020). "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." NeurIPS — source-grounded, explainable AI.

31. Barredo Arrieta, A., et al. (2020). "Explainable Artificial Intelligence (XAI)." Information Fusion, 58.

32. ISO/IEC 42001:2023. Artificial Intelligence — Management System. ISO/IEC.

33. ISO/IEC 23894:2023. Artificial Intelligence — Guidance on Risk Management. ISO/IEC.

Industry & Vendor Quality Solutions

34. Aras Corporation. Aras Innovator — Quality Management System (QMS) solution documentation.

35. Siemens Digital Industries Software. Teamcenter Quality / Closed-Loop Quality documentation.

36. PTC Inc. Windchill Quality Solutions documentation.

37. Dassault Systèmes. 3DEXPERIENCE Quality Management documentation.

38. Autodesk, Inc. Fusion Manage / Quality Management (closed-loop QMS) product documentation.

For more information, contact askus@infosys.com

Infosys®
Navigate your next