FROM DESIGN TO DEPLOYMENT: A DIGITAL ARCHITECTURE STRATEGY

ABSTRACT

The life sciences industry invests heavily in digital technologies, yet pharmaceutical and biopharmaceutical manufacturers continue to struggle with fragmented data landscapes, manual reconciliation, and limited return on digital initiatives. This paper argues that the root cause is architectural, not technological: digital systems designed to optimize individual lifecycle phases fail to preserve the relationships between manufacturing intent, execution reality, and quality evidence as products move from early clinical development through commercial production, technology transfer, contract manufacturing, and product lifecycle management.

This paper introduces the Intent-Driven Digital Core (IDDC), a four-layer architectural framework — Semantic, Event, Context, and Evidence — that preserves manufacturing intent as a persistent, machine-readable construct across the full product lifecycle, from the earliest process development activities through commercial-scale operations and beyond. The IDDC builds on established standards (ISA-88, ISA-95, ICH quality guidelines) and extends them by explicitly addressing semantic persistence, contextual binding at the point of execution, and progressive evidence accumulation. The framework is evaluated through alignment with industry consortium work (Pistoia Alliance CMC Process Ontology, BioPhorum Data Enablement Program), a practitioner case study in alarm management for cell therapy manufacturing, and regulatory requirements including 21 CFR Part 11, EU GMP Annex 11, and current data integrity guidance.

The paper demonstrates how intent-driven architecture accelerates the journey from discovery to patient — technology transfer becomes repeatable rather than bespoke, batch disposition shifts from retrospective reconstruction to continuous evaluation, and AI deployment gains the contextual grounding required for trust in regulated environments. The implications extend beyond technical architecture to organizational governance — treating manufacturing intent as a lifecycle asset requiring explicit stewardship reshapes roles, responsibilities, and operating models.

Limitations, including the framework's design science methodology and scope boundaries, are discussed alongside directions for future empirical validation and cross-industry applicability.

EXECUTIVE SUMMARY

The life sciences industry is undergoing a fundamental shift in how therapies are designed, developed, transferred, and manufactured. Scientific innovation is accelerating, manufacturing modalities are diversifying, and regulatory expectations increasingly emphasize data integrity, traceability, and lifecycle control (FDA, 2011; ICH, 2009). In response, organizations are investing heavily in digital technologies across discovery, development, and manufacturing. Yet despite this investment, many struggle to achieve continuity, scalability, and sustained value from their digital initiatives. While the integration of artificial intelligence (AI) is transforming drug discovery—dramatically increasing both speed and precision (Jiménez-Luna et al., 2021)—development and manufacturing processes have not progressed at a comparable rate, creating structural bottlenecks that delay the delivery of timely and affordable treatments (McKinsey & Company, 2016; ISPE, 2023).

Industry research documents that 70% or more of digital transformation initiatives fail to achieve their stated objectives, with regulated manufacturing environments experiencing particularly high failure rates attributed to underestimation of data context, governance, and semantic consistency

requirements when scaling beyond isolated use cases (McKinsey & Company, 2016; World Economic Forum, 2022; BioPhorum, 2023). This paper argues that architectural discontinuities, rather than individual technologies, are the primary barrier to sustainable digital value. To overcome these barriers, organizations must design and govern a digital architecture and operating model that ensures manufacturing intent, context, and evidence persist from discovery through development, initial design, deployment, routine manufacture, and transfer to new sites. When digital systems are implemented in isolation, critical meaning is lost at lifecycle handoffs (see Figure 1 and Figure 3) requiring repeated interpretation, manual reconciliation, and risk-based judgment under uncertainty.

A robust digital architecture must therefore do more than integrate systems or move data. It must enable consistent interpretation of scientific and manufacturing intent as that intent evolves from early development through commercial production, preserving the relationships between process design, execution reality, and quality expectations. This continuity is essential to supporting reliable scale-up, effective technology transfer (see Figure 6) as defined in WHO Technical Report Series guidelines (WHO, 2011; 2022), and confident batch disposition (see Figure 5) in increasingly complex manufacturing networks.

The strategy described in this paper is grounded in an intent-driven operational model, in which manufacturing intent is explicitly defined, governed, and instantiated across lifecycle phases. This model, consistent with ICH Q8 quality-by-design principles (ICH, 2009) and ICH Q10 pharmaceutical quality system requirements (ICH, 2008), facilitates tighter alignment between discovery and development activities—where process understanding is generated—and manufacturing execution—where that understanding must be applied consistently under regulated conditions. By treating data, context, and evidence as lifecycle assets rather than phase-specific artifacts, the approach reduces friction at transitions and enables earlier, more informed decision-making in accordance with ICH Q12 lifecycle management principles (ICH, 2019).

At the heart of this strategy is the **Intent-Driven Digital Core (IDDC)**: a conceptual framework introduced in this paper to describe a persistent architectural foundation that preserves semantic meaning, execution context, and quality-relevant evidence across the end-to-end lifecycle. Unlike conventional data warehouses, point-solution integrations, or technology stacks, the Intent-Driven Digital Core does not replace existing systems or prescribe specific technologies. Instead, it provides a unifying architectural construct that allows diverse systems to operate against a shared understanding of meaning, time, and state—ensuring data is not only transmitted, but interpretable and actionable throughout the lifecycle in compliance with ALCOA+ (Attributable, Legible, Contemporaneous, Original, Accurate, plus Complete, Consistent, Enduring, and Available) principles (FDA, 2018; MHRA, 2018). These same characteristics become prerequisites for trustworthy AI: as organizations adopt machine learning and large language models, the attributability, provenance, and contextual completeness that ALCOA+ demands grow more essential, not less, because models inherit the integrity of the data and context they consume. This foundation enables downstream capabilities to be deployed with greater trust, scalability, and regulatory confidence.

The Intent-Driven Digital Core is composed of interdependent capabilities, including explicit semantic definitions grounded in industry ontologies (Pistoia Alliance, 2023; 2024; NIIMBL, 2021), event-driven representations of execution reality aligned with ISA-88 batch control concepts (IEC, 2013), contextual binding of process, equipment, material, and quality information consistent with ISA-95 integration models (IEC, 2013), and structured evidence suitable for quality review and regulatory compliance. Together, these capabilities enable organizations to move from ad-hoc digital initiatives toward an integrated, governed environment that supports traceability, compliance, and innovation. Industry research on digital convergence in manufacturing and life

sciences similarly highlights the importance of architectural foundations and governance in scaling digital value beyond isolated projects (World Economic Forum, 2022; ISPE, 2023; BioPhorum, 2023).

Importantly, architecture is not merely a technical exercise; it is fundamentally socio-technical. Preserving and operationalizing manufacturing intent across the lifecycle requires redefining roles and responsibilities, including ownership of intent, stewardship of shared data assets, integration of quality evidence, and leadership across digital tech transfer activities—organizational capabilities emphasized in the ISPE Pharma 4.0 Operating Model (ISPE, 2019). Without enabling people to operate within the architectural construct, even well-designed technical solutions fail to deliver sustained value.

This paper develops these concepts across the manufacturing lifecycle, demonstrating how an intent-driven digital architecture improves continuity, reduces manual reconciliation, enhances trust in analytics and AI in accordance with emerging AI governance frameworks (NIST, 2023; ISO/IEC, 2023), and supports scalable technology transfer. Drawing on industry research and practical manufacturing realities documented through consortium participation (Pistoia Alliance CMC Process Ontology project; BioPhorum Data Enablement Program), it presents a strategy that places architectural continuity at the center of digital execution in regulated life sciences environments.

Key Takeaways

Digital architecture determines value realization. Sustainable digital transformation in life sciences manufacturing depends on architectures that preserve intent and context across the full product lifecycle.

Intent-Driven Digital Core as the strategic foundation. An intent-centric core that harmonizes semantics, execution events, context, and evidence enables trustworthy operations, analytics, AI, and scale.

People and roles are part of the architecture. Organizational capabilities—not just tools—must evolve to steward intent and evidence across functional boundaries.

Outcomes over components. Architectural continuity enables scalable manufacturing, quality, and technology transfer rather than incremental, phase-specific improvements.

CONTENTS

1. INTRODUCTION: THE ARCHITECTURE GAP IN DIGITAL MANUFACTURING

The life sciences industry is experiencing sustained pressure to accelerate the translation of scientific innovation into reliable, scalable, and compliant manufacturing operations. Advances in discovery science, including high-throughput experimentation, computational biology, and AI-enabled drug design, are generating new therapeutic candidates at unprecedented speed. At the same time, manufacturing organizations are required to support increasingly diverse modalities, compressed development timelines, and expanding global production networks—all while maintaining stringent regulatory compliance and product quality standards.

In response, companies have invested heavily in digital technologies across the lifecycle, including electronic laboratory systems (LIMS/LES), manufacturing execution systems (MES), advanced analytics platforms, digital twins, and artificial intelligence. These investments are often justified by their potential to improve efficiency, enhance insight, and enable faster decision-making (Westerman et al., 2014; Vial, 2019). However, despite widespread adoption, many organizations struggle to translate digital capability into sustained, end-to-end value across discovery, development, and manufacturing. Research on digital transformation consistently demonstrates that technology deployment alone does not guarantee business value—successful transformation requires simultaneous development of both digital capability and organizational leadership capability (Westerman et al., 2014).

Industry analyses show that digital initiatives frequently stall beyond pilot phases, fail to scale across sites, or require significant manual intervention to remain operationally viable. This pattern is particularly evident at lifecycle transition points—such as process scale-up, site deployment, and technology transfer—where meaning, context, and accountability must be preserved across organizational and system boundaries. The result is a landscape characterized by fragmented data, duplicated effort, and increasing reliance on expert interpretation to reconcile inconsistencies between systems and stages.

This paper argues that these challenges are not primarily technological in nature. Rather, they reflect a fundamental architecture gap between how manufacturing knowledge is generated, represented, and governed during discovery and development, and how it is executed, reviewed, and transferred during routine manufacturing operations. Enterprise architecture theory has long recognized that information systems architecture must provide logical constructs for defining and controlling the integration of system components (Zachman, 1987), and that architecture serves as strategy by enabling consistent execution across organizational boundaries (Ross et al., 2006). When digital systems are implemented to optimize individual lifecycle phases or functional domains in isolation, they fail to preserve the relationships between scientific intent, execution reality, and quality evidence—a pattern that integration architecture literature identifies as semantic discontinuity (Hohpe & Woolf, 2003). Over time, this fragmentation manifests as rework, delay, and increased operational risk.

The consequences of this gap are most visible in areas that demand lifecycle continuity. Technology transfer requires process understanding, control strategies, and quality expectations to be reinterpreted and re-encoded for each new site or partner. Batch disposition depends on reconstructing execution context and evidence after the fact, often across multiple systems. Advanced analytics and AI struggle to gain trust when data lacks consistent semantics, provenance, or temporal alignment. In each case, the organization compensates for architectural shortcomings through manual review, local workarounds, and tacit knowledge held by a limited number of experts.

Recent industry research highlights this structural challenge. Studies on digital transformation in regulated manufacturing environments note that organizations frequently underestimate the importance of data context, governance, and semantic consistency when scaling digital solutions beyond isolated use cases (McKinsey & Company, 2016; Deloitte, 2021). Similarly, pharmaceutical-specific maturity frameworks such as the ISPE Pharma 4.0 Digital Maturity Model (ISPE, 2023) and the BioPhorum Digital Plant Maturity Model (BioPhorum, 2023) emphasize that digital value is constrained not by analytics capability, but by the ability to integrate process knowledge, execution data, and quality evidence into a coherent, lifecycle-aware foundation. These industry consensus models identify architectural integration and semantic interoperability as foundational maturity dimensions that must precede advanced analytics deployment.

Addressing this architecture gap requires a shift in perspective. Rather than viewing digital transformation as a sequence of technology deployments, organizations must design for lifecycle continuity from the outset—what digital transformation research characterizes as moving from disruption response to strategic value creation (Vial, 2019). This involves explicitly defining manufacturing intent, ensuring that intent can be instantiated consistently across systems, and preserving the contextual and evidentiary relationships necessary for quality, compliance, and learning. Architecture, in this sense, becomes the mechanism by which discovery, development, and manufacturing are unified—through shared semantic foundations rather than uniform tooling. Industry reference frameworks including TOGAF (The Open Group, 2018) and domain-specific maturity models (ISPE, 2023; BioPhorum, 2023) similarly emphasize that sustainable digital value requires architectural foundations that enable semantic consistency across organizational boundaries.

The remainder of this paper develops an architectural strategy to address this gap. It introduces manufacturing intent as the foundational construct, defines the Intent-Driven Digital Core (IDDC) as the mechanism for preserving intent and context across the lifecycle, and explores the implications for people, governance, technology transfer, batch disposition, and advanced analytics. Together, these elements form a cohesive approach to moving from fragmented digital initiatives toward sustainable digital continuity in life sciences manufacturing.

1.1. Key Terms Introduced in This Paper

The following terms are introduced or used with specific architectural intent in this paper. Definitions are provided upfront to ensure consistent interpretation across lifecycle phases, functions, and disciplines.

1.1.1. CORE ARCHITECTURAL CONCEPTS

Term	Definition
Digital Continuity	The operational outcome achieved when manufacturing intent, execution context, and quality evidence are preserved across discovery, development, manufacturing, technology transfer, and product lifecycle management including post-approval changes and product retirement without re-interpretation or manual reconstruction. Digital continuity is the goal; the IDDC is the architectural mechanism for achieving it.
Intent-Driven Architecture	An architectural approach in which manufacturing intent is explicitly defined, governed, and bound to execution and evidence across the lifecycle, rather than being implicitly embedded in systems, documents, or local practices.
Intent-Driven Digital Core	A persistent architectural construct that preserves semantic meaning, execution events, contextual relationships, and quality evidence across the

Term	Definition
(IDDC)	manufacturing lifecycle, enabling downstream capabilities without loss of intent or meaning. The IDDC is the primary artifact introduced by this research.
Architectural Continuity	The condition in which manufacturing meaning, intent, and evidence remain interpretable and traceable despite changes in systems, sites, products, or lifecycle phase.
Architectural Discontinuity	The structural condition in which manufacturing meaning, context, or evidence is lost, reinterpreted, or reconstructed at lifecycle transitions due to inadequate architectural provisions for semantic and contextual persistence. Architectural discontinuity is the central problem this paper addresses; the IDDC is designed to prevent it.
Semantic Interoperability	The capability to exchange data between systems such that the meaning of that data is preserved and consistently interpreted by all parties, beyond mere syntactic or format compatibility. Semantic interoperability requires shared ontologies or formal vocabularies and is a prerequisite for lifecycle continuity.
Architectural Construct	A conceptual structure that defines responsibilities, relationships, and integration patterns independent of specific platforms, tools, or implementations. The IDDC is an architectural construct; specific technology choices (e.g., event streaming platforms, ontology tooling) are implementation decisions.

1.1.2. CORE CAPABILITIES OF THE IDDC

Term	Definition
Semantic Layer	The architectural capability responsible for encoding shared definitions, relationships, and meaning for processes, materials, equipment roles, and quality attributes.
Event Layer	The capability that captures manufacturing execution as time-coherent events and state transitions, enabling accurate reconstruction of what occurred, when it occurred, and under what conditions.
Context Layer	The capability that binds execution events to relevant process, material, equipment, site, and quality context at the point of generation, eliminating downstream inference.
Evidence Layer	The capability that accumulates execution data as review-ready evidence, preserving lineage, provenance, integrity, and suitability for quality and regulatory decision-making.

1.1.3. REGULATORY AND QUALITY CONCEPTS

Term	Definition
Control Strategy	A planned set of controls, derived from current product and process understanding, that ensures process performance and product quality (ICH Q10, 2008). Control strategies span the product lifecycle and must evolve with enhanced process understanding.

Process Understanding	A comprehensive body of knowledge regarding a process and its relationship to product quality, including identification of critical process parameters, their acceptable ranges, and the scientific rationale for those ranges (ICH Q8(R2), 2009). Process understanding is generated during development and must be preserved and accessible throughout the commercial lifecycle.
Quality by Design (QbD)	A systematic approach to pharmaceutical development that begins with predefined objectives and emphasizes product and process understanding and process control, based on sound science and quality risk management (ICH Q8(R2), 2009). QbD principles inform how manufacturing intent should be defined and governed.
Design Space	The multidimensional combination and interaction of input variables and process parameters that have been demonstrated to provide assurance of quality (ICH Q8(R2), 2009). Design space represents formalized process understanding that must be preserved across lifecycle transitions.

294 1.1.4.LIFECYCLE AND OPERATING MODEL CONCEPTS

Term	Definition
Manufacturing Intent	The explicit articulation of process objectives, constraints, control strategies, and quality expectations that define how a product is intended to be manufactured, independent of any specific execution system or site.
Lifecycle Asset	Any architectural element (such as intent definitions, semantic models, or evidence structures) that must persist and evolve across multiple lifecycle phases rather than being created and discarded within a single phase.
Capability Stewardship	Ongoing accountability for the health, evolution, and fitness-for-purpose of architectural capabilities, distinct from time-bounded project delivery.
Intent Ownership	Explicit accountability for defining, maintaining, and governing manufacturing intent across lifecycle phases to ensure consistent interpretation and controlled evolution.

295 1.1.5.OUTCOMES ENABLED BY THE ARCHITECTURE

Term	Definition
Continuous Batch Disposition	A quality operating model in which batch release readiness is evaluated progressively as evidence is generated, rather than reconstructed retrospectively at the end of execution.
Digital Technology Transfer	The ability to transfer manufacturing processes between sites or partners using portable intent, comparable execution context, and reusable evidence rather than site-specific re-interpretation.
Downstream Capabilities	Digital capabilities such as batch disposition, analytics, AI, and digital technology transfer that consume architectural foundations but do not define or own core meaning.

296 1.1.6.AI-RELEVANT ARCHITECTURAL CONCEPTS

Term	Definition

AI Enablement (Architectural Sense)	The condition in which data, context, and evidence are sufficiently explicit, stable, and governed to support scalable, trustworthy AI deployment in regulated manufacturing environments.
Semantic Drift	The gradual loss or divergence of meaning that occurs when definitions or assumptions change without controlled governance, often reintroduced through manual workarounds.
Concept Drift	The degradation of AI model performance caused by changes in underlying processes, data distributions, or context that are not reflected in model governance or retraining.
Socio-Technical Architecture	An architectural perspective that treats people, roles, and decision rights as integral components of the system, alongside technology and data.

2. MANUFACTURING INTENT AS THE ARCHITECTURAL ANCHOR

A recurring challenge in life sciences manufacturing is not the absence of process knowledge, but the inability to preserve that knowledge coherently as it moves across lifecycle phases. The ICH quality guidelines establish process understanding—comprehensive knowledge of the process and its relationship to product quality—as foundational to pharmaceutical development (ICH, 2009). During discovery and development, this understanding is generated through experimentation, modeling, and iterative learning consistent with Quality by Design (QbD) principles (Yu, 2008; Rathore & Winkle, 2009). As products progress toward commercialization, process understanding must be translated into executable instructions, control strategies, and quality expectations that can be applied consistently across sites and over time (ICH, 2008). In practice, this translation is rarely seamless, creating what FDA guidance characterizes as gaps between process design, qualification, and continued verification stages (FDA, 2011).

This paper defines manufacturing intent as the structured expression of what must be done, why it must be done, under what constraints, and how success is determined in a manufacturing context. This concept synthesizes and extends established regulatory constructs: ICH Q8's design space (the multidimensional combination of input variables and process parameters demonstrated to provide quality assurance), ICH Q10's control strategy (planned controls derived from process understanding), and ICH Q8's critical quality attributes and critical process parameters (ICH, 2009; ICH, 2008). Manufacturing intent encompasses process objectives, critical parameters and ranges, material and equipment roles, quality attributes, and acceptance criteria. Importantly, it also includes the rationale behind these elements—the relationships and assumptions that explain why particular controls or decisions matter, what Garcia-Munoz et al. (2010) characterize as the "knowledge component" of design space that is often lost in translation to manufacturing.

Manufacturing intent is therefore broader than a recipe, batch record, or specification. It is the unifying construct that links scientific understanding to operational execution and quality oversight, fulfilling what ICH Q10 envisions as lifecycle knowledge management (ICH, 2008). When intent is not explicitly defined and governed, it becomes fragmented across documents, systems, and teams, increasing the likelihood of divergence as products scale and manufacturing networks expand.

2.1. Fragmentation of Intent Across the Lifecycle

In many organizations, manufacturing intent is implicitly encoded multiple times as a product progresses from development to routine manufacture. Early-stage intent — from the earliest clinical manufacturing at laboratory and pilot scale through process development and scale-up —

may be captured in development reports, experimental data sets, and design rationales. As processes are industrialized, portions of that intent are re-expressed in process descriptions, MES configurations, automation logic, and quality documentation. Each translation introduces the risk of semantic drift, simplification, or loss of context.

Organizations that defer intent capture to later lifecycle stages — treating it as documentation to be assembled rather than architecture to be built — pay for that deferral at every subsequent transition. The cost is not merely inefficiency; it is structural. Without an explicit intent record from the earliest process development activities, each downstream team must reconstruct rationale through conversations, cross-departmental inquiries, and institutional memory — resources that degrade as personnel change and timelines compress. The architectural imperative is clear: the intent record must begin when the first process design decision is made, and it must persist and accumulate as the product matures. Equally important, that record must remain accessible to every downstream team that depends on it; an intent record that exists but cannot be readily found or queried imposes the same reconstruction cost as no record at all.

This fragmentation is particularly pronounced at lifecycle transition points, such as process scale-up, site deployment, and technology transfer. At these stages, intent is often reconstructed from artifacts rather than transferred as a coherent whole. The result is increased reliance on expert interpretation, local workarounds, and retrospective justification to reconcile differences between design assumptions and execution reality.

Industry research and regulatory observations have documented this challenge extensively. Studies on pharmaceutical manufacturing digitalization note that organizations frequently struggle to maintain alignment between process knowledge generated during development and the controls applied during commercial production (McKinsey & Company, 2016; BioPhorum, 2023). This gap is particularly pronounced in complex or rapidly evolving modalities such as biologics and cell therapies, where process-product interdependencies demand tighter knowledge continuity (Nasr et al., 2017). FDA has similarly observed that lifecycle stages often operate as disconnected activities rather than integrated elements of a unified quality system (FDA, 2011). Without a consistent architectural anchor, intent becomes contextualized to a specific phase or system rather than preserved across the lifecycle—what Kourti (2006) describes as the "informatics gap" between process development and manufacturing. Figure 1 illustrates this contrast between traditional discontinuity patterns and the continuity enabled by an intent-driven approach.

2.2. Intent as a Lifecycle Asset

Treating manufacturing intent as an architectural anchor requires a shift in how it is managed, aligning with ICH Q12's vision of product lifecycle management through established conditions and post-approval change management protocols (ICH, 2019). Rather than being viewed as a collection of phase-specific artifacts, intent must be treated as a lifecycle asset—explicitly defined, versioned, and governed as it evolves. This approach operationalizes ICH Q10's concept of a pharmaceutical quality system that spans the product lifecycle and enables continual improvement through knowledge management (ICH, 2008). Intent must adapt as scientific understanding improves and operational constraints change; however, those changes should be deliberate, traceable, and interpretable across all consuming systems and stakeholders—characteristics that ICH Q9(R1) identifies as essential for knowledge-based quality risk management (ICH, 2023).

In practice, intent does not evolve as a simple linear progression. The product lifecycle moves forward — from early development through clinical manufacturing, scale-up, technology transfer, and commercial production — but at specific points, new evidence or operational experience triggers a structured loopback to an earlier design decision. A process parameter established during development may require revision when commercial-scale equipment introduces different mixing dynamics. A control strategy designed for a single clinical site may need adaptation when transferred to a contract manufacturer with different infrastructure. Each loopback is targeted and additive: the original intent remains as the anchor, the trigger for change is documented, the revised approach is justified by evidence that the intended outcome is preserved, and the updated rationale feeds forward into all subsequent lifecycle activities. This pattern — forward progression with structured, versioned loopbacks — ensures that the intent record compounds in depth and defensibility as the product matures, rather than degrading at each transition. For regulators and auditors, the result is a traceable narrative: not merely what was done, but why it was done, what changed, and why the change is justified.

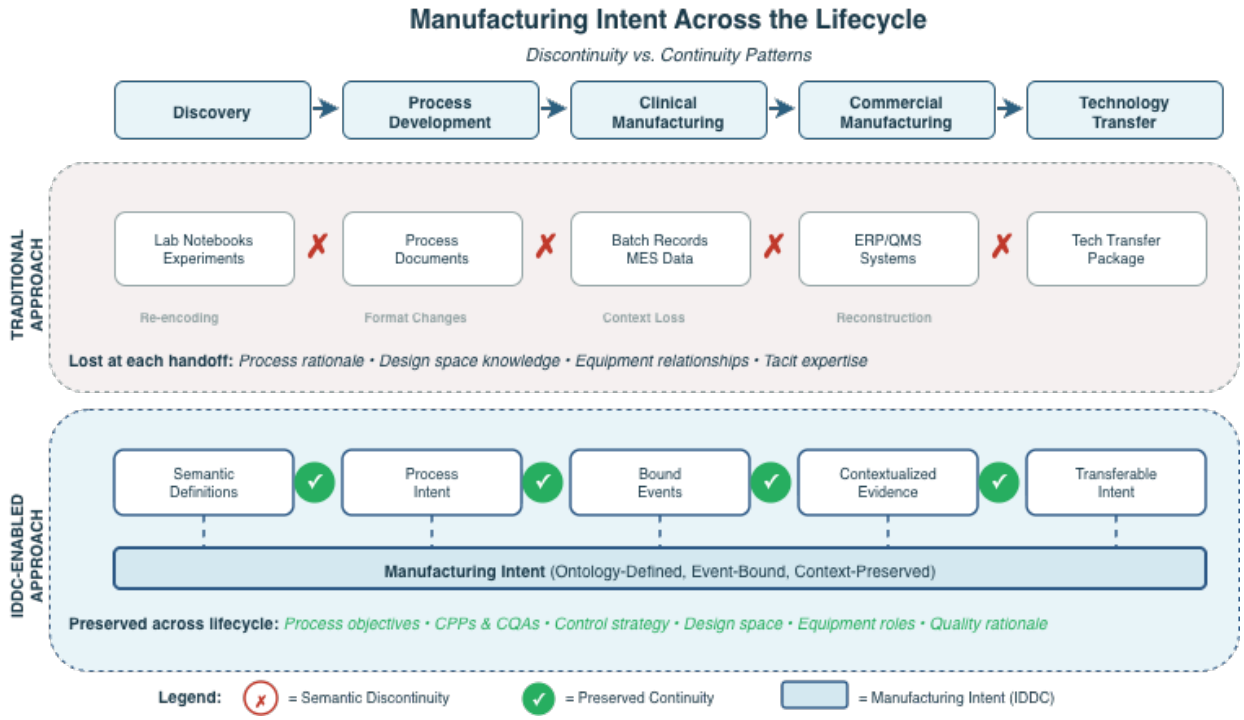

Figure 1. Manufacturing Intent Across the Lifecycle: Discontinuity vs. Continuity Patterns.
The IDDC preserves semantic meaning across lifecycle transitions through ontology-defined intent and context binding.

From an architectural perspective, this means that intent must be:

- Explicit, rather than embedded implicitly in system configurations or documents
- Structured, such that relationships between process elements, constraints, and quality expectations are machine- and human-interpretable
- Persistent, remaining accessible and meaningful across lifecycle phases and sites

By anchoring the architecture around intent, downstream systems are no longer responsible for inferring meaning from isolated data points. Instead, they instantiate and consume a shared representation of what the organization intends to achieve and control through manufacturing.

2.3. Relationship Between Intent, Execution, and Evidence

Manufacturing intent sits at the intersection of design, execution, and quality. It defines what execution should accomplish, but it must also be reconciled continuously with what occurs on the shop floor. Execution events, in turn, generate evidence that demonstrates whether intent has been met.

When intent is weakly defined or inconsistently applied, this relationship breaks down. Execution data becomes difficult to interpret, and quality review shifts toward reconstructing context rather than assessing conformance. Conversely, when intent is explicit and architecturally anchored, execution events can be contextualized at the point of generation, and evidence can be accumulated progressively rather than assembled retrospectively.

This relationship is central to modern quality paradigms, which emphasize lifecycle knowledge management and continuous verification rather than episodic review. FDA's Process Analytical Technology (PAT) framework established the principle that quality should be built into products through design and continuous process understanding rather than tested into finished products (FDA, 2004). ICH Q10 extends this vision by requiring pharmaceutical quality systems to integrate process performance and product quality monitoring with corrective and preventive action and knowledge management (ICH, 2008). More recently, regulatory discussions of real-time release testing and continuous manufacturing have reinforced that quality oversight should be integrated with—not separated from—process execution (Nasr et al., 2017). An intent-driven architecture provides the structural means to achieve this integration without increasing manual burden, enabling what Peterson (2008) characterizes as science-based, rather than specification-based, quality assurance.

2.4. Implications for Architecture Design

Positioning manufacturing intent as the architectural anchor has direct implications for digital architecture design. Architectures optimized solely for data movement or system interoperability are insufficient; they must also support the definition, instantiation, and governance of intent across systems and lifecycle phases (as illustrated in Figure 2).

This requires architectural constructs capable of preserving semantic meaning, binding intent to execution events, and maintaining traceability between design assumptions and operational outcomes. ISA-88 batch control standards provide foundational concepts for structuring manufacturing procedures and equipment relationships (IEC, 2010), while ISA-95 enterprise-control integration models define how manufacturing operations interface with business systems (IEC, 2013). However, neither standard fully addresses the lifecycle persistence of manufacturing intent or the semantic continuity required across development-to- manufacturing transitions. These requirements motivate the need for a persistent architectural foundation—introduced in the next section as the IDDC—that can sustain intent continuity as manufacturing environments evolve, building upon but extending established automation standards.

By elevating manufacturing intent from an implicit assumption to an explicit architectural construct, organizations establish a stable reference point for digital execution, quality assurance, and lifecycle learning. This anchor enables the unification of discovery, development, and manufacturing not through uniform tooling, but through a shared and governed understanding of purpose and control.

3. THE INTENT-DRIVEN DIGITAL CORE: ARCHITECTURE FOR LIFECYCLE CONTINUITY

Having established manufacturing intent as the architectural anchor, the next challenge is preserving that intent as manufacturing systems, sites, and lifecycle phases change. In most digital manufacturing environments, intent is translated repeatedly—into system configurations, data models, and procedural artifacts, each time risking semantic drift or loss of context. Over time, this fragmentation erodes trust in digital systems and increases dependence on manual interpretation. A central limitation of many digital manufacturing initiatives is the implicit assumption that data integration alone is sufficient. Enterprise integration research has long distinguished between syntactic interoperability (the ability to exchange data) and semantic interoperability (the ability to exchange meaning)—with the latter requiring shared conceptual models rather than merely compatible formats (Hohpe & Woolf, 2003; NIST, 2004). In practice, the challenge in manufacturing is not the movement of data, but the persistence of meaning as systems, sites, and lifecycle phases change.

To address this challenge, this paper introduces the concept of the Intent-Driven Digital Core (IDDC): a persistent architectural foundation designed to preserve manufacturing intent, execution context, and quality-relevant evidence across the full product lifecycle.

The IDDC is not a platform, application, or data repository. Rather, it is an architectural construct that governs how meaning is defined, how execution reality is captured, and how evidence is accumulated and interpreted as manufacturing activities progress from discovery through development, deployment, routine operation, and technology transfer.

3.1. Why "Intent-Driven" and Why "Core"

The Intent-Driven Digital Core is intent-driven because manufacturing outcomes depend fundamentally on preserving the design intent that defines process success. This principle aligns with the digital thread concept articulated by NIST and aerospace manufacturing research: an authoritative, extensible data construct that links product definition, production, and support data across the lifecycle (Hedberg et al., 2017). Scientific understanding, regulatory expectations, and operational experience generate intent that must persist across lifecycle phases. An architecture that fails to explicitly capture and govern this intent forces downstream systems to infer meaning—introducing risk at every handoff, what Grieves and Vickers (2017) characterize as information mirroring failure in digital twin architectures.

The architecture is designed to absorb change while preserving semantic continuity. When intent evolves—whether due to improved scientific understanding, operational learning, or site-specific constraints—that evolution is reflected explicitly and traceably, rather than implicitly through disconnected system updates. The IDDC is dynamic in its responsiveness to manufacturing reality, but anchored by intent.

It is a core because it underpins all downstream digital capabilities. Advanced analytics, artificial intelligence, digital twins, batch disposition, and technology transfer do not extend the IDDC; they depend on it. This positioning reflects established enterprise architecture principles: Ross et al. (2006) demonstrate that sustainable digital capability requires a foundation for execution—a stable architectural base that enables rather than constrains business agility. Without such a core that preserves intent, downstream capabilities operate on reconstructed or inferred context, limiting scalability, trust, and regulatory confidence.

3.2. Architectural Scope and Boundaries

The IDDC occupies the space between manufacturing intent and operational execution, complementing established automation standards. ISA-88 (IEC 61512) provides hierarchical models for batch control—procedural elements, equipment entities, and recipe structures—that govern how manufacturing processes are executed (IEC, 2010). ISA-95 (IEC 62264) defines enterprise-control integration, specifying how manufacturing operations interface with business systems through standardized object models (IEC, 2013). The IDDC does not replace these standards or the control systems, manufacturing execution systems, laboratory systems, or quality platforms that implement them. Instead, it provides the semantic and contextual substrate that allows these systems to interoperate without loss of meaning—addressing lifecycle continuity gaps that ISA-88 and ISA-95 were not designed to span.

Its scope includes:

- The formal representation of manufacturing intent and its evolution
- The binding of intent to execution events as they occur
- The preservation of context necessary to interpret execution data
- The accumulation of evidence suitable for quality review, release, and learning

Equally important are its boundaries. The IDDC deliberately avoids:

- Workflow orchestration logic
- Visualization or analytics presentation layers
- Tool-specific integration patterns

By maintaining these boundaries, the core remains durable even as individual systems are replaced, upgraded, or reconfigured.

3.3. Constituent Capabilities of the Intent-Driven Digital Core

The IDDC comprises four interdependent architectural layers—Semantic, Event, Context, and Evidence—that together ensure lifecycle continuity (see Figure 2). These layers represent architectural capabilities rather than discrete technology products; a single platform implementation might span multiple layers, while distributed implementations might deploy each layer using specialized tools. What remains constant across implementations is the responsibility structure: the Semantic Layer must provide shared definitions, the Event Layer must capture time-coherent execution reality, the Context Layer must bind events to their manufacturing context, and

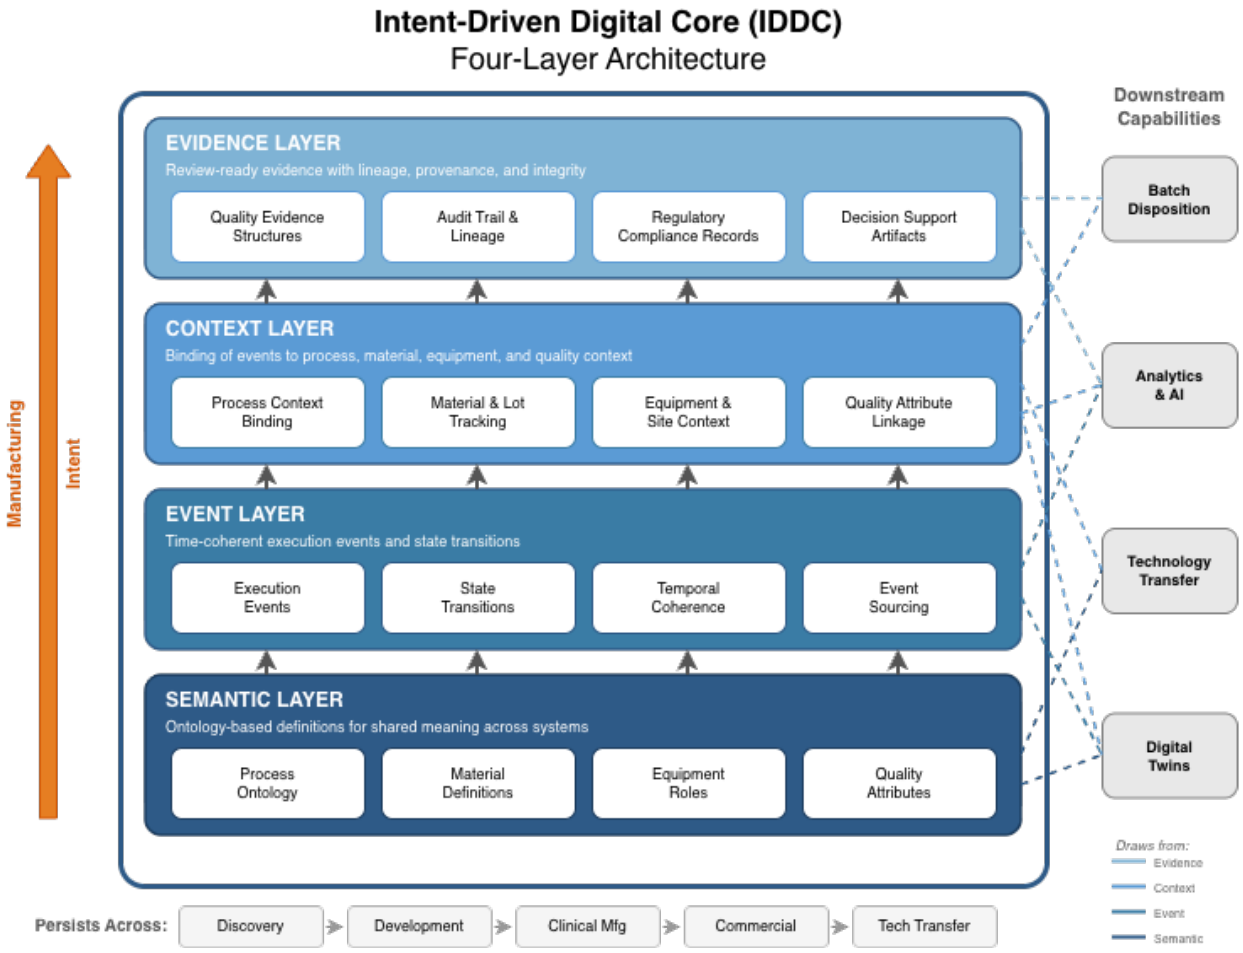

the Evidence Layer must accumulate review-ready quality evidence.

Figure 2. The Intent-Driven Digital Core: Four-Layer Architecture enabling lifecycle continuity through semantic definition, event capture, context binding, and evidence accumulation.

As illustrated in Figure 2, these four layers work in concert: the Semantic Layer provides the vocabulary and relationships that give meaning to manufacturing concepts; the Event Layer captures what actually happens during execution; the Context Layer binds those events to their relevant process, material, and equipment context; and the Evidence Layer structures this contextualized execution data as reviewable evidence. This layered approach ensures that downstream consumers—whether quality reviewers, analytics systems, or receiving sites during

technology transfer—can interpret manufacturing data without reconstructing meaning from fragments.

3.3.1. SEMANTIC CAPABILITY

At its foundation, the core provides a shared semantic representation of manufacturing intent through formal ontological structures. An ontology, in the information systems sense, is an explicit specification of a conceptualization—a formal representation of entities, relationships, and constraints within a domain (Gruber, 1993). This includes explicit definitions of processes, materials, equipment roles, quality attributes, and their relationships. Industry initiatives such as the Pistoia Alliance CMC Process Ontology and NIIMBL Biopharmaceutical Manufacturing Ontology demonstrate growing consensus that such formal semantic foundations are essential for manufacturing interoperability (Pistoia Alliance, 2023; NIIMBL, 2021). These definitions ensure that systems and stakeholders interpret manufacturing meaning consistently, regardless of where or how execution occurs. (see Figure 4 for the broader ontology landscape)

3.3.2. EVENT CAPABILITY

Manufacturing execution is inherently event-driven. Fowler (2017) distinguishes three meanings of "event-driven": event notification, event sourcing, and event-based state transfer—each with distinct architectural implications. The IDDC employs event sourcing principles: rather than storing only current state, it captures execution reality as time-aligned events that constitute an append-only log of state transitions (Hohpe, 2006). This preserves the sequence, timing, and conditions under which actions occur, enabling accurate reconstruction of execution history without retrospective inference—a capability essential for regulatory compliance and root-cause analysis.

3.3.3. CONTEXT CAPABILITY

Events alone are insufficient without context. The ALCOA+ principles—attributable, legible, contemporaneous, original, accurate, complete, consistent, enduring, and available (FDA, 2018)—codify what regulators expect from manufacturing records, and each principle depends on context being bound to data at the point of generation rather than inferred afterward. ISA-95 defines object models for the key manufacturing entities—personnel, equipment, physical assets, and material—that provide the contextual dimensions against which execution must be understood (IEC, 2013). The IDDC extends this approach by binding execution events to relevant process, material, equipment, and —where applicable—patient or lot context at the point of generation. This binding occurs contemporaneously with execution, satisfying the ALCOA+ principle of contemporaneous recording (FDA, 2018). Context binding eliminates the need for downstream systems to infer context after the fact—a common source of data integrity findings in regulatory inspections.

3.3.4. EVIDENCE CAPABILITY

Quality and compliance depend on evidence, not data volume. The core structures execution outputs as evidence, preserving lineage, provenance, and integrity. This evidence is accumulated continuously and remains directly consumable by quality processes, including batch disposition and regulatory review.

Together, these capabilities enable the IDDC to function as a living representation of manufacturing reality, rather than a retrospective record assembled from disparate sources.

3.4. Lifecycle Persistence and Change

A defining characteristic of the IDDC is its ability to persist across lifecycle phases without semantic degradation—realizing the digital thread vision of authoritative data continuity from design through production and support (Hedberg et al., 2017). Manufacturing intent defined during

development must remain interpretable during site deployment, routine execution, and technology transfer. When changes occur, they are incorporated into the core in a controlled manner that preserves traceability between prior and updated intent, consistent with ICH Q12 principles of established conditions and post-approval change management (ICH, 2019).

This persistence supports learning across the lifecycle. Execution outcomes can be evaluated against the intent in effect at the time of manufacture, enabling meaningful comparison, root-cause analysis, and continuous improvement without re-interpretation of historical context— what the NIST smart manufacturing program characterizes as model-based manufacturing intelligence (NIST, 2016).

3.5. Relationship to Downstream Capabilities

The IDDC does not deliver value in isolation. Its purpose is to enable other capabilities to function reliably and at scale.

Batch disposition becomes continuous rather than episodic because evidence is accumulated with context by design (Figure 5). Technology transfer becomes repeatable because intent is explicitly encoded and transferable rather than reconstructed (Figure 6). Advanced analytics and AI become credible because they consume semantically stable, contextualized data rather than fragmented proxies.

In this way, the IDDC acts as the causal mechanism linking architectural design decisions to operational, quality, and analytical outcomes.

> ***Key Takeaway***
>
> The IDDC reframes digital manufacturing architecture around meaning, execution reality, and evidence, rather than data movement alone. By preserving manufacturing intent and context across the lifecycle, it establishes the foundation necessary to unify discovery, development, and manufacturing in a regulated, scalable, and trustworthy manner.

4. LIFECYCLE DISCONTINUITY FROM DESIGN TO DEPLOYMENT

Despite significant advances in digital tooling across the life sciences value chain, most organizations continue to experience pronounced discontinuities as products move from discovery to development and into routine manufacturing. Industry research consistently documents that digital transformation initiatives in pharmaceutical manufacturing fail to achieve stated objectives at rates exceeding 70%, with lifecycle transitions identified as primary failure points (illustrated in Figure 1) (McKinsey & Company, 2016; BioPhorum, 2023). These discontinuities are not limited to system interfaces or data availability; they reflect deeper breaks in how manufacturing intent, context, and accountability are represented and carried forward across lifecycle phases—what knowledge management research characterizes as the challenge of transferring sticky, context-dependent knowledge across organizational boundaries (Argote & Ingram, 2000).

During discovery and early development, scientific and process knowledge is generated in environments optimized for exploration and learning. Experimental data, models, and design rationales are created rapidly, often with a high degree of contextual richness but limited formal structure. As products progress toward commercialization, this knowledge must be translated into

executable processes, control strategies, and quality expectations suitable for regulated manufacturing environments. In practice, this translation is rarely continuous.

4.1. Handoffs as Points of Semantic Loss

Lifecycle transitions such as scale-up, site deployment, and technology transfer are typically treated as handoffs between organizational units, systems, or documentation sets. WHO guidance on technology transfer defines this process as the systematic procedure for transferring documented knowledge and experience during all phases of development and manufacturing, emphasizing that transferred knowledge must be sufficient for the receiving unit to understand and apply the process (WHO, 2011; WHO, 2022). At each handoff, manufacturing intent is partially re-encoded: design assumptions are summarized, control strategies are simplified, and context is abstracted to fit the constraints of downstream systems.

When intent is not explicitly represented as a lifecycle construct, these handoffs become points of semantic loss. Critical relationships between process parameters, material attributes, equipment capabilities, and quality outcomes are weakened or omitted—precisely the process understanding that Kogut and Zander (1992) characterize as difficult to articulate and transfer because it is embedded in organizational routines. Downstream teams are then forced to reconstruct meaning from incomplete artifacts, relying on expert interpretation, informal communication, and retrospective justification to bridge the gaps. Figure 3 visualizes this semantic loss pattern across typical lifecycle handoffs.

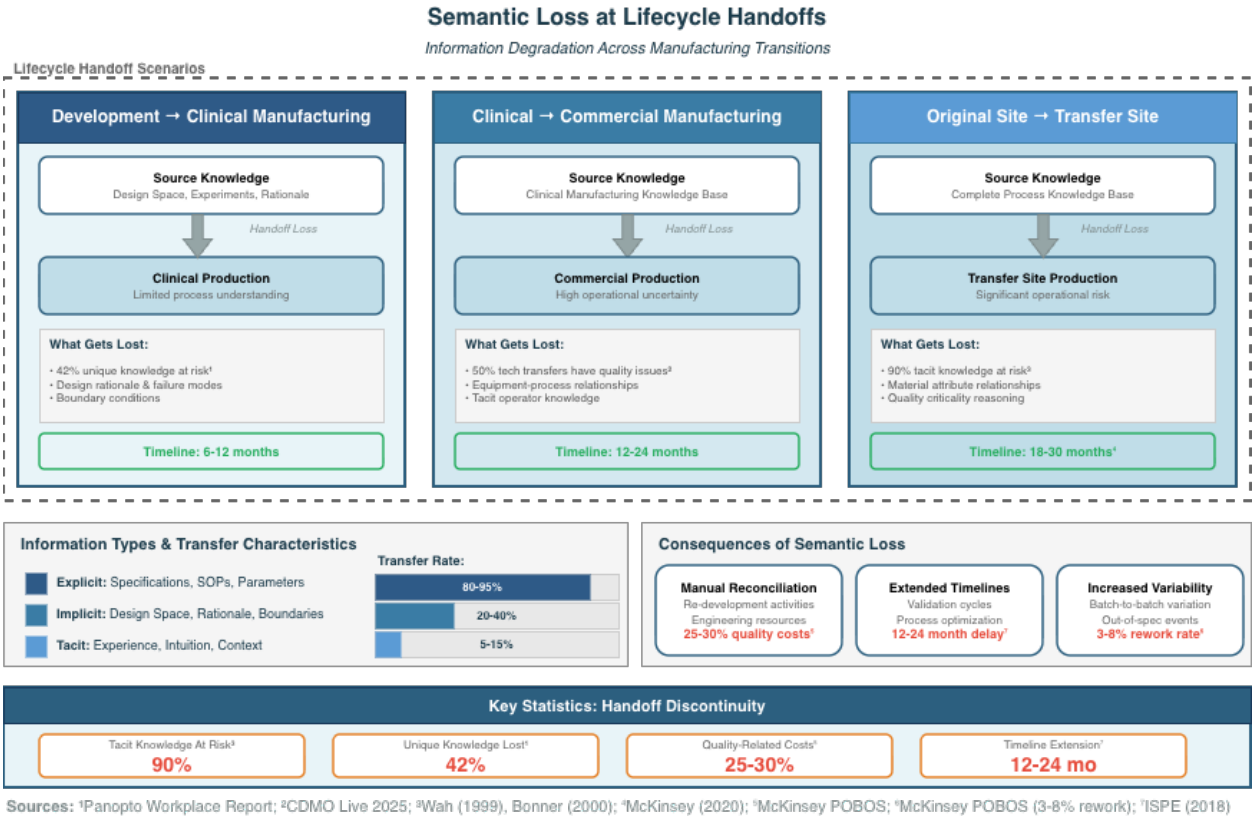

Figure 3. Semantic Loss at Lifecycle Handoffs: Knowledge types, transfer rates, and organizational consequences.

Figure 3 highlights three critical handoff scenarios—development to clinical, clinical to commercial, and original site to transfer site—where semantic loss characteristically occurs. The figure distinguishes between explicit information (specifications, SOPs, batch records) that typically transfers successfully, implicit knowledge (design space, control strategy rationale) that is often weakened, and tacit expertise (troubleshooting experience, process intuition) that rarely survives transfer. Industry data confirms the consequences: knowledge gaps—rather than purely technical issues—are widely identified as a leading cause of extended transfer timelines. The IDDC addresses this by making implicit knowledge explicit through ontological representation and by capturing tacit expertise as contextualized execution patterns.

This pattern is especially evident in multi-site and external manufacturing networks, where differences in equipment, systems, and local practices amplify the effects of incomplete or ambiguous intent transfer. ISPE technology transfer guidance documents that receiving sites routinely require 12-24 months to achieve comparable process performance, with knowledge gaps and context loss identified as primary contributors to extended timelines (ISPE, 2018). The result is increased variability, longer startup timelines, and greater reliance on tacit knowledge held by a small number of individuals—a fragile dependency that BioPhorum identifies as a critical barrier to manufacturing network scalability (BioPhorum, 2020).

> The author's experience confirms this pattern. When introducing a cell therapy manufacturing process to SCADA-augmented production, the team required nine months from initial engagement to detailed design. This included over one hundred hours of requirements workshops, a three-day intensive cross-functional session to define the process flow at major unit operation level, and two additional weeks of detailed design interviews. The discovery was humbling: beyond reading through batch records, there was limited structured understanding of how the process actually functioned. Multiple clinical process versions existed—with no mapping of how one version had derived from another or why changes had been made. Even after implementation partners produced painstakingly detailed graphical process flow diagrams, workshop participants continued to identify inaccuracies, local variations, and anticipated future changes. The knowledge existed, but it was distributed across individuals, embedded in practice, and never captured in a form that could be transferred without extensive person-to-person reconstruction. A well-constructed, intent-driven digital architecture would have substantially reduced this discovery burden—not by eliminating workshops entirely, but by bringing structured process knowledge, version lineage, and design rationale to inform the conversation from the start.

A related pattern is observable when products launch into manufacturing without a structured intent record. In one observed instance involving multiple concurrent product launches at a clinical-stage manufacturer, process characterization records were incomplete, control strategies were undocumented, and the rationale for key equipment selections and process design decisions existed only in the knowledge of individual subject matter experts distributed across departments. When those experts were unavailable, manufacturing and engineering teams were forced to reconstruct intent through ad hoc cross-departmental conversations — a process that introduced delays, risked misinterpretation, and could not withstand regulatory scrutiny. When process changes were needed to accommodate clinical or commercial realities, the absence of a documented intent record meant that each change required de novo justification rather than

traceable departure from an established baseline. An intent-driven architecture, had it been in place from early development, would have provided the structured rationale, the evidence of original design decisions, and the traceability to demonstrate that changes preserved the intended outcome — exactly the defensibility that investigators and auditors require under ICH Q10 and Q12.

4.2. Fragmentation Across Systems and Functions

Lifecycle discontinuity is further reinforced by functional and system silos. Discovery, development, manufacturing, and quality organizations often operate distinct digital ecosystems, each optimized for local objectives and constraints. While integration efforts may connect these systems at a technical level, they rarely ensure continuity of meaning.

As a consequence, the same manufacturing concepts such as a critical process parameter or quality attribute may be represented differently across systems, with no authoritative source of truth. ISA-95 (IEC 62264) defines standard object models for manufacturing operations, yet implementations frequently diverge in how these models are instantiated locally, creating semantic heterogeneity despite nominal standards compliance (IEC, 2013). Temporal misalignment between data sources compounds the problem, making it difficult to establish a coherent view of execution reality. Over time, digital systems become sources of data, but not of shared understanding—a condition that NIST semantic interoperability research identifies as the core barrier to manufacturing system integration (NIST, 2006).

This fragmentation shifts effort away from value-adding activities toward reconciliation. Engineers, operators, and quality reviewers spend significant time assembling, validating, and interpreting information rather than evaluating performance or managing risk. While this manual effort often succeeds in the short term, it does not scale as product portfolios and manufacturing networks grow.

4.3. Consequences for Quality, Transfer, and Insight

The impact of lifecycle discontinuity is most visible in areas that depend on continuity by design.

In batch disposition, evidence required for release is frequently assembled after execution from multiple systems, each with its own assumptions and context. Review effort increases as reviewers must reconstruct what occurred, under what conditions, and whether intent was met. This episodic approach places a heavy cognitive burden on quality organizations and limits opportunities for earlier risk identification.

In technology transfer, the absence of digitally transferable intent forces each receiving site to reinterpret process understanding and quality expectations. ICH Q5E establishes that manufacturing process changes—including site transfers—require demonstration of comparability to ensure product quality, safety, and efficacy are maintained (ICH, 2004). Even when documentation is complete, differences in execution context and system representation require local adaptation, extending timelines and introducing variability. Scale-up and transfer research documents that process understanding losses during technology transfer contribute significantly to manufacturing deviations in the first 12-18 months of commercial production (Farid, 2007; Garcia-Aponte et al., 2017).

In advanced analytics and AI, fragmented semantics and inconsistent context undermine trust. Models trained on data that lacks stable meaning or provenance struggle to generalize across sites or time, reinforcing skepticism and limiting adoption beyond exploratory use cases.

Across these domains and others, organizations compensate for architectural gaps through increased oversight, manual controls, and localized expertise. While effective in regulated

environments, these compensatory mechanisms constrain scalability and slow the realization of digital value.

4.4. Architectural Roots of Discontinuity

Crucially, these challenges do not arise from poor execution or insufficient digital investment. They stem from architectures that were not designed to preserve intent, context, and evidence across the lifecycle. Enterprise architecture research demonstrates that sustainable digital capability requires designing for integration from the outset, treating data and meaning as shared enterprise assets rather than application-specific resources (Ross et al., 2006). When digital systems are implemented as phase-specific solutions optimized for local objectives, continuity becomes an operational burden rather than an architectural property.

Without a persistent architectural construct to anchor manufacturing intent and bind it to execution reality, continuity must be recreated repeatedly through people and process. This approach is inherently fragile and increasingly unsustainable as manufacturing complexity grows.

Addressing lifecycle discontinuity therefore requires more than improved integration or documentation. It requires an architectural foundation capable of maintaining semantic continuity as products, sites, and systems evolve. The following sections examine how such an architecture is enabled—through explicit semantic constructs, socio-technical alignment, and lifecycle-aware governance—and how its effectiveness can be evaluated through outcomes such as batch disposition, technology transfer, and advanced analytics.

Key Takeaway

Lifecycle discontinuity manifests not as missing data but as lost meaning—semantic loss at handoffs, fragmented context across systems, and reconstructed (rather than transferred) process understanding. These patterns are architectural in origin and cannot be resolved through improved documentation or integration alone.

5. ONTOLOGY AS A CONTRACT OF MEANING

The lifecycle discontinuities described in the previous section arise not from a lack of data, but from the absence of a shared, durable understanding of what that data represents. When manufacturing systems exchange information without a common semantic foundation, continuity of intent cannot be assumed. Instead, meaning must be inferred, reconciled, or reinterpreted (often manually) at each lifecycle transition.

To address this challenge, digital architectures must incorporate an explicit mechanism for defining and governing meaning. In this context, ontology is not an academic abstraction, but a practical architectural construct: a contract of meaning that ensures manufacturing intent, execution context, and quality evidence are interpreted consistently across systems, sites, and lifecycle phases.

5.1. Ontology as an Architectural, Not Academic, Construct

In manufacturing environments, the term “ontology” is often associated with knowledge graphs or formal logic models. While these techniques can be valuable, the architectural role of ontology is more fundamental. An ontology defines the entities, roles, relationships, and constraints that give data its meaning within a specific domain.

In an intent-driven digital architecture, ontology serves as the authoritative reference for:

- What constitutes a process, material, equipment role, or quality attribute
- How these elements relate to one another
- How changes to one element affect interpretation of others

This semantic contract allows systems to exchange information without embedding meaning directly into point-to-point integrations or local data models. Instead, meaning is externalized, governed, and reusable.

5.2. From Data Integration to Semantic Interoperability

Traditional system integration focuses on syntactic compatibility, ensuring that data can be transmitted and stored. While necessary, this approach does not guarantee that data will be interpreted consistently. As a result, the same concept may appear under different names, structures, or assumptions across systems, reinforcing the fragmentation described in Section 4.

Ontology enables semantic interoperability by decoupling meaning from system implementation. When systems reference shared semantic definitions, they no longer need to encode local interpretations of manufacturing concepts. This is particularly important in regulated environments, where consistency of interpretation underpins quality, compliance, and trust.

This shift from data integration toward semantic interoperability is already reflected in emerging industry initiatives. The Pistoia Alliance CMC Process Ontology extends established manufacturing frameworks (specifically ISA-88 and ISA-95) by defining shared semantic representations of processes, parameters, and manufacturing stages across chemical and biologic modalities.

NIIMBL, in collaboration with the OAGi, has released biopharmaceutical manufacturing ontologies to harmonize representations of concepts across the value chain. These life-science-specific efforts align with broader industrial ontology work led by NIST through initiatives such as the Industrial Ontologies Foundry (IOF), which provide reference models and design patterns for manufacturing-domain semantic frameworks. Collectively, these initiatives reflect an industry-wide recognition that scalable interoperability depends on shared meaning, not merely connected systems. Figure 4 illustrates this ontology landscape, showing how foundational standards build toward domain-specific applications.

Established manufacturing frameworks, such as ISA-95 and related object models, provide a useful starting point by defining common abstractions for equipment, materials, and processes. However, these frameworks must often be extended to reflect modality-specific requirements, complex material genealogy, and evolving process understanding. Ontology provides the mechanism to formalize and govern these extensions without fragmenting the underlying architecture.

5.3. From Semantic Definition to Operational Deployment

The preceding discussion establishes what ontology provides architecturally—a governed, externalized contract of meaning that enables semantic interoperability across systems and lifecycle phases. A complementary question, less frequently addressed in the literature, concerns what ontology demands of the systems and organizations that deploy it: where does it reside, how do operational systems consume it, and through what mechanisms do semantic definitions translate into executable, traceable manufacturing operations?

A proposed deployment architecture addressing these questions is developed in Appendix B, using cell and gene therapy (CGT) manufacturing as a bounded scope. The architecture describes six layers through which ontological relationships are realized in practice: governed ontology definition, ontology access and distribution infrastructure, MES instantiation aligned with ISA-95 structural objects, integration-layer contextualization, data platform materialization, and consumption by quality and analytical processes. Of particular practical significance is the access layer—the mechanism by which consuming systems reference, bind to, and stay synchronized with the governed ontology—which addresses the fundamental deployment question of where the ontology resides and how systems ensure they operate against the current semantic version. The central architectural principle is that ontology functions as a governed semantic backbone—systems map to it rather than execute it—with ISA-95 providing structural execution models and ontology supplying the meaning, roles, and traceability that make execution data interpretable across boundaries.

CGT manufacturing was selected as the stress-test context because it intensifies every dimension of the deployment problem. Autologous therapies demand mandatory lot genealogy, non-collapsible lineage, patient linkage as a first-class semantic construct, and disposition semantics that apply to a specific patient's therapy rather than a statistical population of product units. The ontology deployment architecture that satisfies these requirements inherently accommodates less demanding modalities—establishing CGT as a rigorous validation boundary for the IDDC's Semantic Layer.

It should be noted that this deployment architecture represents a design science artifact: a proposed solution validated by alignment with international standards (ISA-95, ISA-88, BFO, IOF Core) and consortium consensus (Pistoia Alliance, NIIMBL), rather than a validated implementation. The biopharmaceutical industry does not yet have a published example of end-to-end ontology deployment in manufacturing. The contribution is the architecture itself—providing an engineering blueprint grounded in the ontology landscape described above and the standards ecosystem that governs it.

5.4. Role of Ontology in Preserving Manufacturing Intent

Manufacturing intent, as defined in Section 2, cannot persist across the lifecycle unless its meaning is explicit and interpretable. Ontology enables this by formally encoding the relationships that define intent; for example, how a critical process parameter relates to a quality attribute, or how an equipment role constrains execution conditions.

When intent is instantiated in execution systems, ontology ensures that execution events are interpreted in the correct semantic context (the Semantic Layer shown in Figure 2). This eliminates the need for downstream systems or reviewers to infer intent from isolated data points. Instead, execution data is inherently meaningful because it is generated and consumed within a shared semantic framework.

Internal and external ontology initiatives provide the semantic foundation required to preserve manufacturing intent across lifecycle phases, enabling design assumptions, execution reality, and quality evidence to remain coherently linked as processes scale, transfer, and evolve.

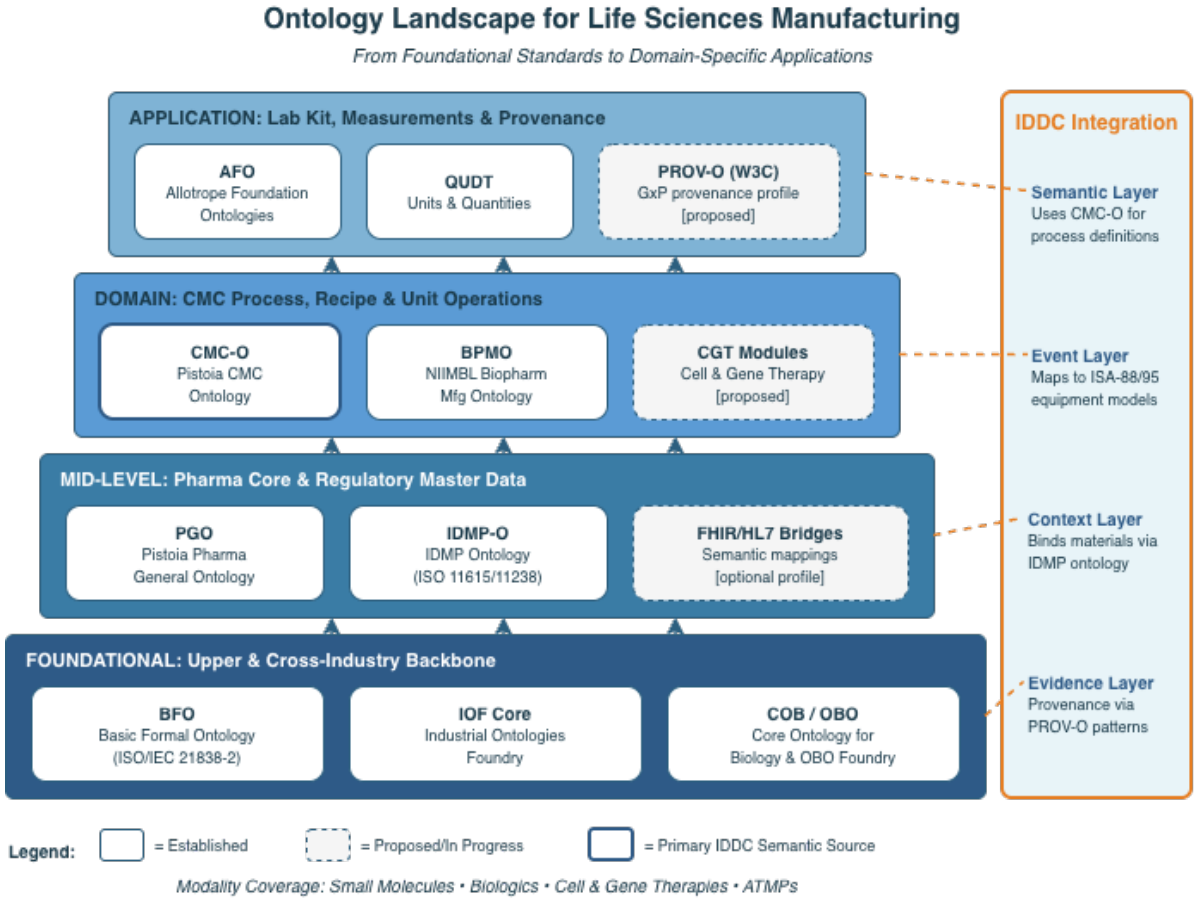

Figure 4. Ontology Landscape: Foundational to Domain-Specific for Life Sciences Manufacturing. The CMC Process Ontology (CMC-O) provides the primary semantic foundation for the IDDC Semantic Layer.

5.5. Ontology as an Enabler of Evidence and Trust

Quality decisions depend not only on data availability, but on confidence in how that data was generated and what it represents. Ontology contributes directly to this confidence by providing traceable, governed definitions of meaning that persist over time.

When execution data and contextual information are linked through shared semantic constructs, evidence accrues as execution unfolds rather than being reconstructed after the fact. Reviewers can evaluate outcomes against the intent in effect at the time of manufacture, with clear lineage between design assumptions, execution events, and quality conclusions.

This semantic stability also underpins trust in advanced analytics and AI. Models trained on data that is semantically consistent and contextually bound are more likely to generalize across sites and time, reducing skepticism and enabling broader adoption.

5.6. Governance of Meaning Across the Lifecycle

Because ontology defines meaning, it must be governed with the same rigor applied to other lifecycle assets. Changes to semantic definitions, whether driven by new process understanding, operational learning, or regulatory feedback, must be deliberate, versioned, and traceable.

This governance does not imply centralized control of all systems, but it does require clear ownership of semantic constructs and disciplined change management. Without such governance, ontologies risk fragmenting in the same way as ungoverned data models, undermining their architectural value.

When properly governed, ontology becomes the connective tissue between intent, execution, and evidence—enabling the IDDC to function as designed.

> ***Key Takeaway***
>
> Ontology provides the semantic contract that allows manufacturing intent to be interpreted consistently across systems, sites, and lifecycle phases. By externalizing and governing meaning, ontology transforms digital integration from data exchange into true lifecycle interoperability—enabling continuity, trust, and scalability in regulated manufacturing environments.

6. EVENT, CONTEXT, AND EVIDENCE BY DESIGN

Digital manufacturing architectures often treat execution data as a byproduct of system operation rather than as a primary architectural concern. Data is captured because systems generate it, not because it has been intentionally designed to serve downstream interpretation, quality decisions, or lifecycle learning. As a result, execution records are frequently incomplete, temporally misaligned, or insufficiently contextualized to support confident decision-making without manual reconstruction. Regulatory expectations for data integrity (FDA, 2018; MHRA, 2018; PIC/S, 2021) increasingly require that manufacturing records be attributable, contemporaneous, and complete at the point of generation—requirements that retrospective assembly inherently struggles to satisfy. These requirements align with the ALCOA+ framework, which has become the de facto standard for data integrity expectations across global regulatory authorities.

An intent-driven digital architecture requires a different approach. Rather than asking how data can be collected and integrated after execution, it must be designed so that events, context, and evidence are bound together at the point of execution, by design.

6.1. Manufacturing Execution as an Event-Driven Reality

Manufacturing is inherently event-driven. Actions occur in time, under specific conditions, involving defined materials, equipment, and process states. Fowler (2017) distinguishes three architectural patterns for event-driven systems—event notification, event-carried state transfer, and event sourcing—each with different implications for how execution history is preserved and replayed. However, many digital systems still represent execution primarily as time-series data streams or aggregated records, detached from the events that give them meaning.

When execution is treated as a sequence of discrete, time-aligned events (rather than as isolated measurements), what occurred, when it occurred, and under what conditions becomes explicit rather than inferred. This event-centric representation enables accurate reconstruction of

execution history without reliance on retrospective interpretation or manual correlation across systems.

Designing for events does not require replacing existing control or execution systems. It requires ensuring that execution data is captured and structured in a way that reflects the actual sequence and state transitions that matter for process understanding and quality assessment.

This distinction does not diminish the value of high-frequency, time-series data at the point of generation. Detailed signal-level data remains essential for real-time monitoring, control, and engineering troubleshooting, where fine-grained temporal resolution enables rapid diagnosis of equipment behavior and process dynamics. However, not all time-series data needs to propagate unchanged through downstream systems. As execution data moves beyond local operational contexts, temporal coherence and contextual relevance become more important than raw resolution. Event-based abstractions allow downstream consumers to interpret execution outcomes accurately without carrying the full burden of signal-level detail that is meaningful primarily at the point of execution. This layered approach aligns with Hohpe and Woolf's (2003) principle that integration architectures should transform data to match consumer needs rather than propagating raw operational detail.

6.2. Context Binding at the Point of Execution

Events alone are insufficient without context. The same action or measurement can have different significance depending on the process step, material state, equipment role, or control strategy in effect at the time. When context is not bound at execution, downstream systems and reviewers must infer it after the fact, often with incomplete information.

An intent-driven architecture binds context to execution events at the point of generation, using the shared semantic constructs established through ontology as shown in the Context Layer of Figure 2. This includes, at minimum:

- Process context (step, phase, or operation)
- Material context (lot, batch, or patient linkage where applicable)
- Equipment context (role, configuration, or capability)
- Quality context (applicable limits, specifications, or control strategies)

By binding this context explicitly, execution data becomes immediately interpretable. Downstream consumers no longer need to reconstruct meaning from timestamps, system identifiers, or procedural assumptions. This design directly supports the expectations of 21 CFR Part 11 and EU GMP Annex 11 (European Commission, 2011), which require that electronic records be complete, accurate, and readily retrievable throughout their retention period.

6.3. From Data Capture to Evidence Generation

Quality and compliance decisions depend on evidence, not raw data. The FDA's process validation guidance (FDA, 2011) and ICH Q10 (ICH, 2008) both emphasize that quality assurance depends on contemporaneous, contextually complete records that demonstrate process understanding and control—precisely the characteristics that distinguish evidence from mere data capture. Evidence is data that has been contextualized, attributed, and structured to support an explicit claim such as conformance to intent, suitability for release, or identification of risk. Evidence generation, as opposed to process data generation (see Figure 5 for how this enables continuous disposition), therefore depends not on preserving all execution data at its highest resolution, but on preserving the right abstractions of execution aligned to manufacturing intent.

In many manufacturing environments, evidence is assembled retrospectively. Execution data is extracted from multiple systems, aligned manually, and reviewed against expectations after the fact. This episodic approach increases review burden and limits opportunities for early detection of deviations or emerging trends.

A contrasting architectural approach is demonstrated by an alarm management platform implemented at a cell therapy manufacturing organization (see Appendix A). Prior to implementation, alarms generated across diverse equipment types accumulated in individual systems, each with its own interface and storage approach. When alarms were logged by operators, they were typically reviewed only when a deviation triggered an investigation—a pattern common across the industry given resource constraints, but one that creates temporal disconnect between alarm occurrence and quality assessment. In autologous therapies, where one batch represents one patient, the consequences of delayed alarm visibility extend beyond operational metrics to direct patient impact.

The platform architecture explicitly addresses the Event, Context, and Evidence layers of the IDDC framework. The Event Layer captures alarms with their equipment, process, and batch relationships. The Context Layer binds human interpretation—documenting causation, treatment, and impact assessment—contemporaneously with execution. The Evidence Layer produces structured output including alarm lineage, audit trail, and the decision framework supporting batch disposition. This progressive evidence accumulation enables the continuous disposition model described in Section 8, while the platform's extensible architecture provides a foundation for the AI-assisted capabilities discussed in Section 10.

Designing for evidence generation changes this dynamic. When events and context are captured together, evidence can be accumulated progressively and continuously as execution occurs. Lineage, provenance, and integrity are preserved by design, reducing reliance on manual reconciliation and enabling earlier, more focused quality oversight. PIC/S guidance on data management and integrity (PIC/S, 2021) explicitly recognizes that data governance embedded in system design is more robust than procedural controls applied after the fact.

6.4. Temporal Coherence and Traceability

A common failure mode in digital manufacturing is temporal incoherence: data streams that cannot be reliably aligned in time or attributed to the correct execution state. This undermines traceability and erodes trust in digital records. EU GMP Annex 11 (European Commission, 2011) requires that audit trails record time-stamped, operator-attributed changes to electronic records—a requirement that presupposes the temporal coherence that many manufacturing data architectures fail to provide.

An event-centric, context-bound architecture enforces temporal coherence by anchoring all execution data to a common event timeline. This allows design intent, execution reality, and quality outcomes to be evaluated in relation to one another without ambiguity.

Traceability, in this model, is not achieved through exhaustive logging alone, but through structured relationships between intent, events, and evidence. This approach supports both regulatory expectations for traceable, attributable records (FDA, 2018; PIC/S, 2021) and operational learning without excessive data overhead.

6.5. Architectural Implications

Designing for event, context, and evidence has several important architectural implications:

- Execution systems must expose meaningful events, not just signals, following event-driven patterns appropriate to the integration context (Fowler, 2017; Hohpe & Woolf, 2003)
- Semantic definitions must be available at execution time, not applied later, consistent with the ontological access architecture described in Section 5
- Data pipelines must preserve event ordering and contextual relationships, ensuring temporal coherence as required by 21 CFR Part 11 and EU GMP Annex 11 (European Commission, 2011)
- Evidence generation must be treated as a first-class architectural outcome

Key Takeaway

Event-driven execution, contextual binding, and evidence generation are not downstream analytics concerns; they are architectural design choices. When these elements are designed together, manufacturing data becomes inherently meaningful, reviewable, and reusable across the lifecycle. When they are not, organizations are forced to compensate through manual interpretation and retrospective reconstruction.

7. PEOPLE AS AN ARCHITECTURAL DEPENDENCY

Digital architecture determines not only how information flows, but also who interprets, owns, and acts on manufacturing intent as it evolves across the lifecycle. As scientific understanding is translated into executable processes and quality decisions, architecture shapes where meaning is preserved, where judgment is applied, and where responsibility resides. As architectures evolve from isolated systems toward intent-driven, lifecycle-aware foundations, the role of people shifts accordingly. In this context, people are not external users of digital systems; they are integral components of the design itself. This perspective aligns with socio-technical systems theory (Trist & Bamforth, 1951), which demonstrates that technical and social subsystems must be jointly optimized—optimizing one at the expense of the other produces suboptimal outcomes regardless of the sophistication of either component.

When architectural design fails to account for human roles, decision rights, and capabilities, continuity cannot be sustained. Meaning is reintroduced informally, intent is reinterpreted locally, and the fragmentation described in earlier sections reappears despite technical integration.

7.1. Why Digital Architecture Redefines Roles

Traditional manufacturing organizations evolved around system and functional boundaries. Automation, manufacturing execution, quality systems, and enterprise platforms were implemented and governed independently, with human roles focused on system ownership and reconciliation. Digital maturity was measured by system availability and compliance rather than by continuity of meaning across the lifecycle. Conway's (1968) observation that organizations design systems mirroring their communication structures explains why: when organizational boundaries align with system boundaries rather than with lifecycle flows, the resulting architectures inevitably fragment meaning at each handoff.

An intent-driven digital architecture fundamentally alters this dynamic in three ways.

First, it drives a shift from system ownership to intent ownership. When manufacturing intent is explicitly defined and preserved across lifecycle phases, accountability naturally moves toward those responsible for defining, maintaining, and governing that intent, rather than those operating individual systems. This represents a shift from what Mintzberg (1979) characterizes as machine

bureaucracy—organized around functional specialization—toward a more integrative structure where coordination mechanisms cut across traditional departmental boundaries.

As a result, human effort shifts decisively from manual reconciliation of fragmented records toward informed judgment under uncertainty, where manufacturing context is preserved and intent is explicit by architectural design.

Second, it collapses traditional IT, OT, and Quality boundaries. Semantic consistency, event coherence, and continuous evidence generation require cross-functional collaboration by design. Architectural separation between these domains becomes a source of risk rather than protection.

Third, it increases reliance on human judgment over manual reconciliation. As execution data becomes contextualized and evidence is accumulated progressively, people spend less time assembling information and more time evaluating uncertainty, risk, and deviation significance. This sets the stage for what Malone (2018) describes as collective intelligence: human-machine collaboration in which AI handles routine pattern recognition and evidence assembly while humans exercise judgment on significance, risk, and response—a partnership that produces outcomes neither could achieve independently.

Architecture, in this sense, does not remove people from the process. It elevates the nature of human contribution and increases impact.

7.2. Capability-Based Roles Enabled by the Architecture

As the IDDC becomes the organizing foundation, certain capabilities become architecturally necessary. These should be understood as functional roles rather than prescriptive job titles, consistent with the capability-based organizational models advocated by the ISPE Pharma 4.0 Operating Model (ISPE, 2019) and Galbraith's (2014) framework for designing organizations around strategic capabilities rather than hierarchical structures.

7.2.1. MANUFACTURING INTENT OWNER

Accountable for defining and stewarding manufacturing intent across lifecycle phases, including process objectives, constraints, and quality expectations. This role ensures semantic continuity from development through routine manufacture and transfer.

7.2.2. DATA PRODUCT OWNER

Responsible for reusable, governed manufacturing data assets that serve multiple consumers, including operations, quality, analytics, and AI. The emphasis is on lifecycle fitness, lineage, and interpretability rather than local optimization.

7.2.3. DIGITAL TECHNOLOGY TRANSFER LEAD

Bridges development, manufacturing, and quality to ensure that intent, context, and evidence are digitally transferable across sites and partners without semantic reinterpretation (see Figure 6).

7.2.4. DATA STEWARD

Ensures that execution evidence generated by the architecture is complete, trustworthy, and suitable for continuous disposition and regulatory review, with explicit focus on provenance and traceability.

7.2.5. MANUFACTURING SYSTEMS ARCHITECT

Designs execution and integration models that preserve intent and context across control systems, MES, and data platforms, ensuring architectural consistency as systems evolve.

These capabilities may be distributed across multiple individuals or embedded within existing roles. What matters is that the capability exists and is explicitly recognized.

The roles described above are intentionally framed as architectural capabilities rather than prescriptive job titles. They align closely with established competency frameworks such as those promoted by the International Association of Automation Engineers (IAAE) and smart manufacturing personas defined by CESMII, while operating at a higher level of abstraction. This positioning allows organizations to map these capabilities onto existing roles and career pathways without mandating structural reorganization, while still ensuring that critical ownership of intent, data, and evidence is explicitly addressed.

7.3. Upskilling Requirements Across the Organization

An intent-driven architecture changes not only roles, but also the skills required to operate effectively within it.

For frontline operators, the emphasis shifts toward:

- Interpreting contextualized digital records rather than transcribing data
- Understanding how execution actions contribute to downstream evidence
- Confident interaction with digital batch records and contextual views

For engineers and technical specialists, required skills increasingly include:

- Ontology-aware process and system design
- Event-driven architectural thinking
- Designing for reuse, transferability, and lifecycle persistence

For quality professionals, the architecture enables and demands a transition toward:

- Continuous disposition models
- Data-driven risk assessment
- Evaluation of uncertainty rather than completeness alone

These changes reinforce that digital transformation is not an add-on training exercise, but a redefinition of professional practice within manufacturing. Kotter (1996) identifies the failure to anchor change in organizational culture as the most common reason transformation initiatives regress; in this context, upskilling must be accompanied by changes in performance expectations, career incentives, and operational workflows to sustain the architectural shift.

7.4. Organizational Implications

Architectural continuity requires organizational alignment. Ownership of intent must be explicit and cross-functional, spanning development, manufacturing, quality, and digital functions. Galbraith (2014) demonstrates that organizations pursuing lateral integration require explicit coordination mechanisms—shared goals, liaison roles, and integrative processes—that go beyond

structural reorganization. Delivery models must shift from project-centric implementations toward capability stewardship, where architectural elements are owned, governed, and evolved over time.

Rather than acting as control gates that slow delivery or solution factories that fragment ownership, Centers of Excellence are most effective when positioned as enablers of architectural coherence, shared standards, and lifecycle learning across the organization (ISPE, 2019).

7.5. Failure Modes When People Are Not Aligned

When organizational roles and skills are not aligned with the digital architecture, predictable failure modes emerge:

- Architectural pathways are bypassed through manual workarounds
- Semantic drift is reintroduced via spreadsheets and local tools
- Batch review effort increases rather than decreases
- Advanced analytics and AI outputs are distrusted or ignored

In these cases, the architecture may exist technically but not operationally. Fragmentation reappears not because the technology failed, but because people were not positioned to operate within it. Industry maturity assessments consistently confirm this pattern: BioPhorum's digital transformation work (BioPhorum, 2023) identifies organizational readiness as a primary differentiator between organizations that scale digital value and those that remain in pilot cycles.

Architecture without people alignment recreates the same fragmentation it seeks to eliminate.

Key Takeaway

Treating people as semantic and decision-making components of the architecture is essential to achieving scalable technology transfer, continuous batch disposition, and credible adoption of advanced analytics and AI.

8. BATCH DISPOSITION AS AN ARCHITECTURAL OUTCOME

Batch disposition has historically been treated as a discrete quality activity performed at the conclusion of manufacturing execution. In this model, data is gathered from multiple systems, contextual gaps are resolved through manual review, and quality decisions are made retrospectively based on the completeness of records rather than the continuous assessment of execution against intent. While effective in regulated environments, this approach places a significant operational burden on quality organizations and becomes increasingly fragile as manufacturing complexity, product portfolios, and network scale grow. The FDA's process validation guidance (FDA, 2011) establishes a lifecycle approach encompassing Stage 1 (Process Design), Stage 2 (Process Qualification), and Stage 3 (Continued Process Verification)—yet the episodic disposition model is fundamentally misaligned with Stage 3's expectation of ongoing, data-driven process assurance.

This paper advances a different perspective: batch disposition is not a workflow to be optimized, but an outcome that emerges from architectural design choices. When manufacturing intent, execution events, and quality evidence are preserved coherently across the lifecycle, disposition shifts from an episodic activity to a continuous, evidence-driven capability. Figure 5 contrasts the

traditional retrospective disposition model with the continuous, evidence-driven approach enabled by the IDDC.

8.1. Limitations of Episodic, Retrospective Disposition

Traditional batch disposition relies on reconstructing execution context after manufacturing has concluded. Evidence is assembled from disparate systems, often with inconsistent semantics and incomplete temporal alignment. Reviewers must infer intent, reconcile discrepancies, and apply expert judgment to determine whether a batch meets release criteria. ICH Q10 (ICH, 2008) envisions batch release as an outcome of effective pharmaceutical quality systems with embedded knowledge management—an aspiration that episodic, reconstruction-dependent disposition cannot fulfill at scale.

As articulated in BioPhorum's Digital Batch Disposition Manifesto (BioPhorum, 2023), this retrospective model scales poorly. Manual reconciliation increases with product complexity, while review timelines lengthen and variability in interpretation grows. Even in highly digitalized environments, the absence of architectural continuity forces quality decisions to depend on human reconstruction of meaning rather than on inherently trustworthy evidence.

These limitations are not primarily procedural. They reflect architectures that were never designed to support continuous disposition. Evidence is generated, but not structured; data is available, but not contextualized; intent exists, but is not explicitly bound to execution.

8.2. Continuous Disposition Enabled by Architecture

An intent-driven architecture fundamentally alters how disposition evidence is created and evaluated. When execution is captured as contextualized events and bound to the manufacturing intent in effect at the time, evidence is accumulated progressively, not assembled after the fact. This architectural approach operationalizes Stage 3 continued process verification (FDA, 2011) as a design property rather than a procedural aspiration.

In this model, each execution event contributes incrementally to the disposition decision. Deviations, excursions, and anomalies are evaluated in context, as they occur, rather than being rediscovered during final review. Quality oversight shifts from exhaustive record assembly toward focused assessment of uncertainty and risk.

The IDDC enables this transition by preserving the semantic relationships between process intent, execution reality, and quality expectations. Evidence generated within this core retains provenance, lineage, and temporal coherence by design, allowing quality decisions to be informed continuously throughout the batch lifecycle. This directly supports ICH Q10's emphasis on knowledge management as a lifecycle enabler of product quality (ICH, 2008) and aligns with ICH Q8(R2)'s vision of control strategies that are understood well enough to be monitored continuously (ICH, 2009).

This architectural approach aligns directly with the principles articulated in the Digital Batch Disposition Manifesto, which emphasizes that digitalization alone does not yield faster or more reliable disposition unless evidence is contextualized and trustworthy at the point of generation.

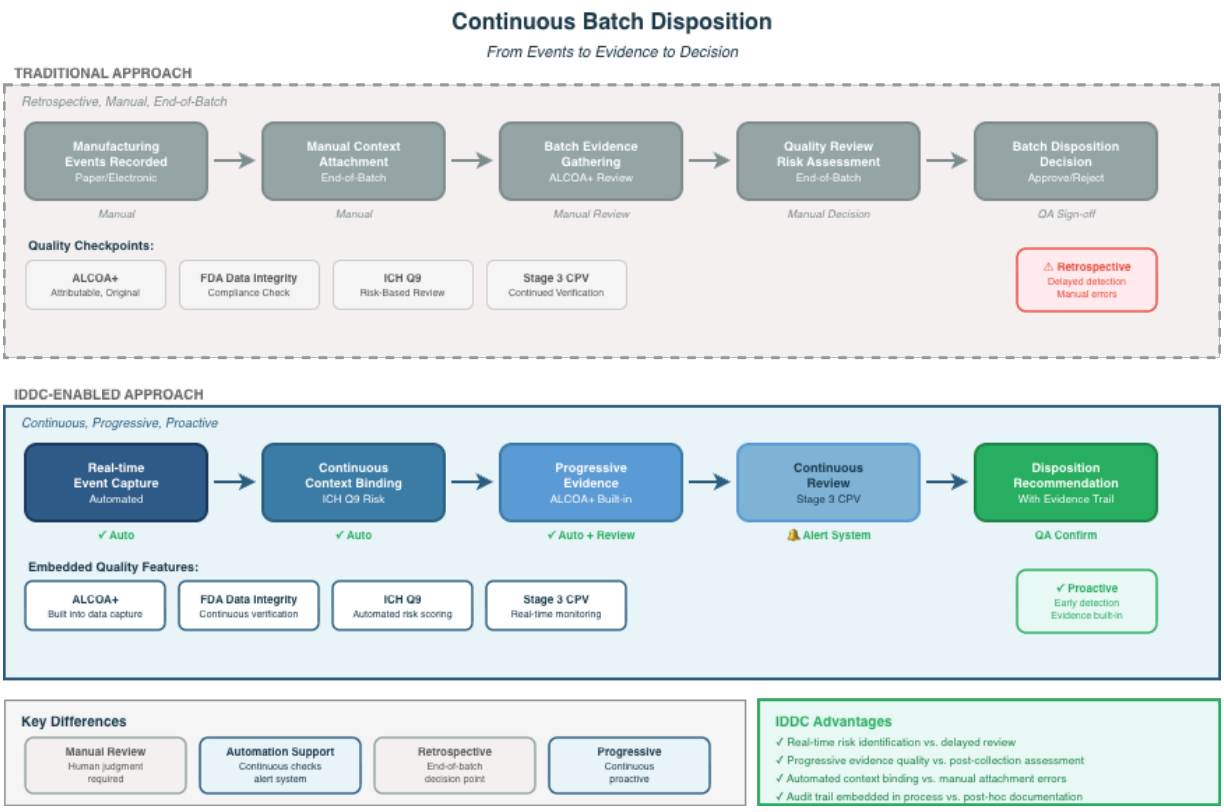

Figure 5. Continuous Batch Disposition: From Events to Evidence to Decision. The IDDC enables progressive evidence accumulation replacing retrospective reconstruction.

Figure 5 illustrates how the IDDC transforms batch disposition from an episodic, retrospective activity into a continuous, proactive capability. In the traditional approach (top), evidence is assembled after manufacturing concludes, requiring manual reconciliation across systems and retrospective context reconstruction. In the IDDC-enabled approach (bottom), manufacturing events are bound to context at the point of execution, evidence accumulates progressively throughout the batch, and quality review shifts from data assembly to exception-focused assessment. This architectural shift enables real-time visibility into batch status, earlier detection of deviations, and disposition decisions that are faster, more consistent, and more defensible—without sacrificing the rigor required for regulatory compliance.

8.3. Evidence as a Lifecycle Asset

Reframing batch disposition as an architectural outcome requires a corresponding reframing of evidence. Evidence must be treated as a lifecycle asset, generated during execution, enriched with context, and preserved for downstream review, learning, and regulatory engagement. The ALCOA+ principles (FDA, 2018) define the attributes that make evidence trustworthy—attributable, legible, contemporaneous, original, and accurate—each of which is more reliably achieved when evidence is accumulated architecturally rather than assembled retrospectively.

When evidence is accumulated continuously, disposition becomes less about verifying record completeness and more about confirming conformance to intent. This distinction is critical. Completeness can be achieved through manual effort; conformance requires architectural alignment.

This approach also enables earlier detection of quality risk. Trends, patterns, and deviations can be assessed against intent while operators are engaged to provide context and corrective action

remains possible, reducing reliance on post hoc justification and minimizing downstream disruption.

8.4. Implications for Quality Organizations

The organizational implications of continuous disposition align with the socio-technical model developed in Section 7: as the architecture assumes responsibility for evidence accumulation, quality professionals redirect their expertise from information assembly to risk evaluation and informed judgment — a transition that strengthens both operational efficiency and regulatory confidence.

This shift is consistent with the socio-technical model described in Section 7. As evidence quality improves, trust in digital records increases, and the credibility of review by exception and other advanced quality models improves accordingly. Importantly, this transition does not require regulatory reinterpretation; it requires architectural readiness. ICH Q12's framework for managing established conditions and post-approval changes (ICH, 2019) provides the regulatory foundation for this evolution: when evidence structures are architecturally governed, the distinction between established conditions and operational parameters becomes tractable rather than ambiguous.

8.5. Disposition as Proof of Architectural Maturity

Batch disposition is one of the most stringent tests of digital manufacturing architecture. Unlike analytics or visualization, disposition decisions carry direct regulatory and patient impact. If an architecture can support timely, confident, and consistent disposition across products and sites, it demonstrates that intent, context, and evidence are truly preserved across the lifecycle. ICH Q9(R1) (ICH, 2023) reinforces that quality risk management must be systematic, science-based, and embedded in the quality system—characteristics that are achievable only when the underlying architecture provides structured, trustworthy evidence for risk assessment.

In this sense, batch disposition functions as a leading indicator of architectural maturity. Architectures that rely on manual reconciliation will struggle to scale disposition, regardless of digital tooling. Architectures that preserve meaning by design enable disposition to compound in efficiency, reliability, and trust as complexity increases.

> ***Key Takeaway***
>
> Batch disposition is not accelerated by digitizing quality workflows, but by architecting for continuous evidence generation and intent preservation. When disposition is treated as an architectural outcome rather than a downstream activity, quality decisions become faster, more consistent, and more scalable without sacrificing regulatory rigor.

9. DIGITAL TECHNOLOGY TRANSFER AS THE ULTIMATE STRESS TEST

Technology transfer is one of the most revealing and least forgiving phases of the manufacturing lifecycle. It requires that scientific understanding, process intent, execution models, and quality expectations remain interpretable as they move across organizational, geographic, and technical boundaries. Unlike steady-state manufacturing, technology transfer exposes whether continuity has been architected or merely assumed.

For this reason, digital technology transfer functions as the ultimate stress test of a digital architecture. Architectures that appear sufficient within a single site or system landscape often fail

under the pressure of transfer, where implicit assumptions, local interpretations, and fragmented representations are forced into the open.

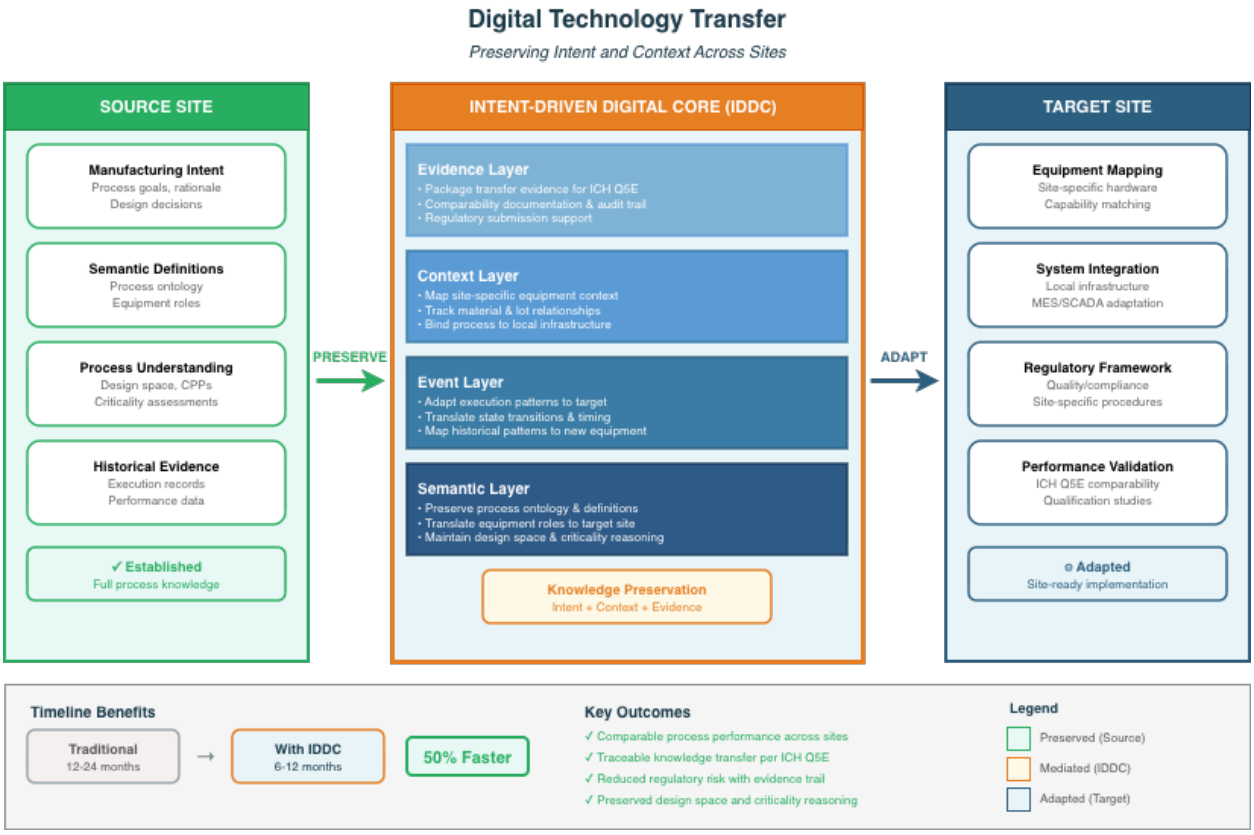

Figure 6. Digital Technology Transfer: Preserving Intent and Context Across Sites. The IDDC is projected to enable significant reduction in transfer timelines while maintaining traceable comparability aligned with ICH Q5E requirements.

As Figure 6 demonstrates, the IDDC proposes a fundamentally different approach to technology transfer. Industry guidance consistently documents that traditional transfer processes require 12–24 months for complex biologics transfers, with some extending beyond 24 months depending on process complexity and receiving site readiness (ISPE, 2018; PDA, 2022; WHO, 2022). Rather than translating process knowledge through documentation that must be reinterpreted at the receiving site, the architecture preserves manufacturing intent in a portable, semantically consistent form. The source site's process definitions, control strategies, and execution patterns are captured within the IDDC's four layers; during transfer, these representations are instantiated at the target site while adapting to local equipment and system configurations. This approach is projected to reduce transfer timelines significantly by eliminating the manual reconstruction of process understanding that WHO TRS 1044 identifies as a primary source of transfer delays (WHO, 2022). Critically, it supports ICH Q5E comparability demonstration through structured, traceable data rather than reconstructed documentation (ICH, 2004), enabling immediate comparability of execution evidence between sending and receiving sites.

9.1. Why Technology Transfer Exposes Architectural Fragility

Technology transfer amplifies complexity along multiple dimensions simultaneously. Processes move from development to manufacturing, from one site to another, or from internal operations to external partners. Equipment, automation platforms, and execution systems change, while regulatory expectations for demonstrated comparability remain constant (ICH, 2004; FDA, 2011).

Under these conditions, any ambiguity in how manufacturing intent is represented becomes a source of delay, rework, or risk.

Historically, organizations have compensated for this fragility through extensive documentation, expert oversight, and extended stabilization periods. While often effective, these approaches scale poorly (ISPE, 2018; WHO, 2022). As product portfolios expand and manufacturing networks become more distributed, reliance on manual interpretation and tacit knowledge becomes a structural constraint that industry guidance increasingly recognizes as a barrier to effective transfer (PDA, 2022).

Many technology transfer challenges attributed to “process complexity” or “site readiness” are in fact architectural in nature. When intent is embedded implicitly in documents, spreadsheets, or system configurations, it must be reconstructed at each receiving site. This reconstruction introduces variability and undermines repeatability, regardless of the maturity of individual systems. ICH Q10 recognizes that effective knowledge management across the product lifecycle is essential for maintaining process understanding during transfer (ICH, 2008), while ICH Q12 provides the regulatory framework for managing post-approval changes that transfers often necessitate (ICH, 2019). Figure 6 illustrates how an intent-driven architecture enables digital technology transfer by preserving meaning across site boundaries.

9.2. Intent Portability as the Core Requirement

Digital technology transfer succeeds only when manufacturing intent is portable. Portability does not imply uniform execution; sites will always differ in equipment, automation, and operational practices. Rather, it requires that the meaning of the process—its objectives, constraints, and quality expectations—remain consistent even as execution models vary. This is the operational expression of what ICH Q5E terms comparability: demonstrating that a product’s quality attributes are not adversely affected by changes in the manufacturing process, including site transfers (ICH, 2004).

An intent-driven architecture enables this portability by decoupling intent from any single system or site implementation. When intent is preserved within the architectural core, receiving sites do not reinterpret the process; they instantiate it. Local adaptation becomes a bounded exercise, governed by explicit constraints rather than implicit assumptions.

This distinction is critical. Translation implies reinterpretation; instantiation implies fidelity. Architectures that support instantiation by design reduce transfer timelines, improve comparability of outcomes, and lower dependence on extended expert involvement. WHO TRS 1044 emphasizes that the completeness and accuracy of transferred process knowledge directly determines transfer outcomes (WHO, 2022); an intent-driven architecture addresses this by making that knowledge structurally complete rather than documentarily complete.

The mechanism through which intent portability is achieved is the ontology access architecture described in this paper’s treatment of ontology as architectural infrastructure. When manufacturing intent is encoded against a governed ontology and systems access that ontology programmatically —rather than through manual transcription into local configurations—the receiving site instantiates the same semantic structure. Process definitions, material specifications, and control strategy parameters retain their governed meaning because both sites resolve them against the same ontological source. This is the architectural equivalent of what ICH Q8(R2) envisions as a design space that is understood well enough to be transferred rather than re-derived (ICH, 2009). The ontology does not eliminate local adaptation—equipment differences, facility constraints, and operational practices will always require site-specific configuration—but it constrains adaptation

to occur within a semantically governed boundary, ensuring that comparability is preserved by design rather than reconstructed through post-transfer analysis.

9.3. Digital Transfer Beyond Documentation

Traditional digital technology transfer efforts often focus on digitizing transfer artifacts: electronic process descriptions, general/master recipes, structured batch records, or standardized parameter lists. While necessary, these artifacts do not guarantee continuity of meaning. As WHO TRS 1044 notes, the transfer of documented knowledge without its underlying context and rationale is a recurring source of transfer failures across the industry (WHO, 2022).

Effective digital transfer requires that execution context and evidence generated at the receiving site can be evaluated against the same intent and quality expectations defined upstream. This demands semantic consistency, event coherence, and contextual binding across environments—capabilities that ISPE's technology transfer guidance identifies as critical success factors for modern transfer programs (ISPE, 2018).

When these conditions are met, early manufacturing runs at a new site generate evidence that is immediately comparable to development and prior production experience. Deviations can be assessed meaningfully, and stabilization efforts can focus on true process differences rather than representational mismatches.

9.4. Transfer Across Networks and Partners

The architectural demands of technology transfer increase further in extended manufacturing networks. Contract development and manufacturing organizations (CDMOs) now execute a significant and growing share of biopharmaceutical production, yet transfer to and between CDMOs remains among the most failure-prone activities in the industry (PDA, 2022). The challenges are both technical and organizational: CDMOs operate their own automation platforms, quality systems, and data architectures, while sponsors retain regulatory responsibility for product quality. Without a shared architectural foundation, transfer relies heavily on contractual artifacts, technology transfer packages assembled as static document sets, and extensive manual oversight—a model that WHO TRS 1044 explicitly identifies as insufficient for complex products (WHO, 2022).

An intent-driven digital architecture provides a mechanism to maintain continuity without imposing uniform systems or tools across partners. By externalizing manufacturing intent and preserving it in a semantically governed, system-agnostic form, sponsors and CDMOs can achieve interoperability while respecting local execution environments (as illustrated in Figure 6). The IDDC's ontology access architecture is particularly relevant here: rather than requiring CDMOs to adopt the sponsor's system landscape, the architecture requires only that both parties resolve process definitions against a shared ontological source. This approach aligns with the ISPE technology transfer guidance's emphasis on transfer protocols that separate process understanding from platform-specific implementation (ISPE, 2018), and with ICH Q10's vision of quality system interoperability across organizational boundaries (ICH, 2008).

This capability is increasingly important as manufacturing strategies emphasize flexibility, resilience, and networked production models. The trend toward platform manufacturing, multi-product facilities, and rapid deployment of new modalities—particularly cell and gene therapies with patient-specific requirements—amplifies the need for architecturally portable intent. Organizations that achieve this portability will be positioned to treat manufacturing network design as a strategic capability rather than a logistical constraint.

9.5. Technology Transfer as Architectural Validation

Technology transfer provides a clear, outcome-based measure of architectural maturity. Architectures that preserve intent, context, and evidence by design enable transfer to compound in efficiency and reliability as experience grows. Architectures that rely on manual reconciliation reset transfer effort with each new site or product—a pattern that current industry guidance consistently identifies as unsustainable at the scale and speed modern manufacturing networks demand (ISPE, 2018; WHO, 2022; PDA, 2022).

In this sense, technology transfer is not merely an operational milestone; it is a validation of whether digital architecture truly supports lifecycle continuity. If intent can be transferred digitally without reinterpretation, the architecture is sound. If not, no amount of tooling or documentation will fully compensate.

> ***Key Takeaway***
>
> Digital technology transfer does not fail because manufacturing systems differ; it fails when manufacturing intent is not portable. Architectures that preserve meaning, context, and evidence across the lifecycle transform technology transfer from a high-risk translation exercise into a repeatable instantiation process, enabling scale, resilience, and confidence in increasingly complex manufacturing networks.

10. AI Enablement as a Consequence of Architecture

Artificial intelligence is increasingly positioned as the next step in life sciences digital transformation. In practice, however, many AI initiatives stall at prototype stages or fail to transition into reliable, regulated operations. The primary reason is rarely algorithmic sophistication. Rather, it is the absence of a lifecycle-ready architectural foundation that preserves meaning, context, and evidence with sufficient integrity for models to be trusted, validated, and sustained over time. Sculley et al. (2015) demonstrate that machine learning systems accumulate "hidden technical debt" through unstable data dependencies, feedback loops, and configuration complexity—debt that compounds rapidly when deployed on architecturally fragmented manufacturing data.

This paper therefore treats AI enablement not as a driver for architecture, but as a consequence of it. AI systems consume manufacturing reality. If manufacturing reality is fragmented, semantically inconsistent, or poorly evidenced, AI inherits those weaknesses, regardless of model performance in controlled settings. Conversely, when manufacturing intent, execution events, and quality-relevant evidence are preserved coherently across the lifecycle, AI becomes more feasible, more scalable, and more credible in regulated environments.

10.1. Why AI Fails to Scale in Regulated Manufacturing

Across the industry, AI proofs of concept frequently succeed in local contexts, then fail to scale across sites, products, or time. Several structural causes recur:

- Context collapse: model features are available, but their meaning varies across systems, sites, or lifecycle phases
- Unverifiable provenance: training and inference data cannot be traced to authoritative sources or aligned to the conditions under which it was generated

- Semantic drift: the same process concept is represented differently across stages, undermining model stability and reproducibility
- Temporal incoherence: event sequences and state transitions cannot be reconstructed reliably, weakening causal inference and operational interpretability
- Credibility gaps: models lack a robust, risk-based justification aligned to a defined context of use, limiting regulatory and business acceptance

Regulators and industry groups have increasingly converged on the idea that credibility depends on defining and bounding context of use, establishing traceable data foundations, and maintaining lifecycle controls on models and their inputs. The FDA's guidance on AI and machine learning in drug and biological product development (FDA, 2023) explicitly frames AI credibility as risk-based and context-dependent, emphasizing the need to justify data relevance and model performance for the intended use.

Similarly, the NIST AI Risk Management Framework (NIST, 2023) emphasizes lifecycle processes, test and evaluation, and governance to manage AI risk, reflecting the broader industry realization that AI is not a one-time build, but a continuously managed capability.

10.2. Intent-Driven Digital Core as the Enabler of AI Readiness

The architectural mechanism that makes AI scalable in this paper's model is the IDDC. By preserving manufacturing intent, semantic meaning, event fidelity, contextual binding, and quality evidence independently of any single system implementation, the IDDC creates the conditions under which AI can be trained and deployed with confidence (see Figure 2).

This is a critical distinction from conventional "AI-ready data platform" narratives. Data platforms can centralize storage and compute, but they do not inherently preserve meaning. AI enablement requires not only data access, but interpretability, traceability, and comparability across the lifecycle.

In practical terms, the IDDC supports AI readiness by ensuring that:

- Intent is explicit and versioned, grounded in the ontology landscape shown in Figure 4, enabling models to learn against authoritative definitions rather than emergent proxies
- Execution is represented as coherent events, supporting temporal reasoning and reducing ambiguity in state-dependent behaviors
- Context is bound at generation, minimizing downstream inference and increasing feature reliability
- Evidence is structured for review, enabling GxP-aligned confidence in model inputs and outputs

- This framing is consistent with the BioPhorum consortium's data characteristics framework (BioPhorum, 2026), which organizes AI-ready data requirements into four groups — data observability, data quality, data availability, and machine understandability — and demonstrates that the characteristics most critical for AI vary by use case archetype rather than being universally applicable. The framework is discussed in detail in the following subsections.

10.3. AI-Ready Data as Lifecycle Evidence, Not Just "Clean Data"

- A common simplification is to equate AI readiness with data quality improvements, such as completeness, accuracy, and standardization. The FAIR principles (Wilkinson et al., 2016)—findable, accessible, interoperable, and reusable—represent an important advance, but even FAIR compliance is insufficient for manufacturing AI: data may be findable and

accessible yet lack the contextual binding and semantic consistency that regulated decision-making demands.

- AI readiness can be distinguished as a multi-factor condition involving:

- data observability (understanding condition, drift, and failure modes)
- metadata versus context (what describes data versus what makes it interpretable)
- lineage and provenance (traceability across transformation and augmentation)
- semantic consistency (stable meaning across systems and phases)
- change dynamics (concept drift, process drift, and evolving intent)

- These distinctions align tightly with what regulators and risk frameworks expect: AI outputs must be interpreted within their context of use and supported by defensible evidence about data fitness, model performance, and lifecycle controls. The EU AI Act (European Union, 2024) classifies AI systems by risk level and imposes stringent requirements on high-risk applications—a category that encompasses many manufacturing quality and safety decisions. ISO/IEC 42001 (ISO/IEC, 2023) further requires organizations to establish AI management systems with explicit governance of data quality, model lifecycle, and risk assessment.

- Within this paper's intent-driven architecture, AI-ready data is best understood as lifecycle evidence that is fit to support a specified analytical purpose. Evidence is not defined by volume or resolution, but by interpretability and traceability. High-resolution time-series signals remain essential for control and troubleshooting at the point of generation, but AI systems operating across lifecycle phases typically require event-centric abstractions and contextual coherence rather than raw signal persistence. This mirrors the architectural rationale established in Section 6.

10.4. Cross-Industry Validation: The BioPhorum Data Characteristics Framework

- The preceding distinction between AI readiness and data quality finds direct validation in the BioPhorum consortium's work on managing structured and unstructured data for AI (BioPhorum, 2026). Developed by a cross-company working group representing major biopharmaceutical manufacturers, the BioPhorum framework establishes that AI-ready data must satisfy characteristics organized into four groups, each answering a distinct question about the data's fitness for AI consumption.

- **Data observability** addresses what can be assessed without domain knowledge — completeness, consistency, timeliness, schema conformance, integrity, lineage, and precision. These are characteristics that data engineering teams can measure and monitor programmatically, providing the baseline assurance that data pipelines are functioning correctly.

- **Data quality** addresses what requires process expertise to evaluate — validity, reasonability, credibility, relevance, representativeness, and authoritativeness. These characteristics depend on understanding how data was generated and what constitutes expected behavior in a specific manufacturing context. They cannot be assessed by schema validation alone.

- **Data availability** addresses what enables AI systems to access data appropriately — findability, accessibility, interoperability, reusability, confidentiality, scalability, and speed of retrieval. These are infrastructure and governance characteristics that determine

whether technically sound, quality-assured data can actually reach the AI systems that need it.

- **Machine understandability** addresses what enables AI systems to interpret data appropriately — contextualization, semantic connection, and clarity. The BioPhorum framework identifies this as the group that represents a genuinely new challenge beyond traditional data management, noting that optimal machine understandability requires alignment between the data's inbuilt characteristics and the organization's ontology, knowledge graph, and process documentation.

- This four-group framework maps directly to the IDDC's layered architecture. Data observability and data quality are primarily served by the Event and Evidence layers — the architectural commitment to capturing execution data with temporal integrity, contextual binding, and traceable provenance. Data availability is served by the integration and data platform layers — the infrastructure that makes semantically grounded data findable, accessible, and interoperable across systems and sites. Machine understandability is served by the Semantic Layer — the ontological foundation that provides AI systems with the structured meaning required to interpret data beyond isolated values.

- The BioPhorum framework further validates a central argument of this paper: AI readiness is not a binary state but a context-dependent assessment. The framework demonstrates that the data characteristics most critical for AI vary by use case archetype. Knowledge and sense-making applications prioritize semantic connection. Classification and detection applications require integrity, precision, and contextual binding. Prediction and forecasting demand temporal completeness and schema conformance. Decision support requires auditability and legal compliance. This archetype-to-characteristics mapping provides an empirically grounded method for determining which architectural capabilities must be most mature for a given AI deployment — directly operationalizing the "context of use" principle discussed in the following subsection.

- The convergence between the BioPhorum framework and the IDDC architecture is not coincidental. Both emerge from the same underlying recognition: that AI in regulated manufacturing fails not because models are inadequate, but because the data they consume lacks the semantic consistency, contextual integrity, and traceable provenance that regulated environments demand. The BioPhorum work validates this at the consortium level, across organizations representing the majority of global biopharmaceutical production.

10.5. From Use Case to Context of Use

- AI initiatives fail when use cases are articulated as "apply AI to X" rather than as defined contexts of use that specify:

- the decision or action the model will inform
- the risk associated with incorrect outputs
- the operational environment in which the model will be used
- the acceptable uncertainty and performance bounds
- the evidence required for trust, review, and validation

- This aligns closely with the FDA's credibility framework (FDA, 2023), which emphasizes defining context of use and establishing credibility proportionate to risk, and with the NIST AI RMF's (NIST, 2023) emphasis on mapping context, measuring performance, and governing lifecycle risks.

- In regulated manufacturing, the implication is direct: AI should be positioned first as a decision-support capability with bounded scope, not as an autonomous controller. Industry frameworks such as ISPE's guidance on AI in pharmaceutical manufacturing (ISPE, 2025) underscore that without lifecycle assurance and validation thinking from the outset, AI programs tend to remain trapped in pilot cycles.

10.6. Architectural Prerequisites for AI Across the Lifecycle

- While AI use cases vary, the architectural prerequisites tend to converge. In the context of this paper's architecture, scalable AI requires the following prerequisites to be satisfied as lifecycle properties, not project deliverables:

10.6.1. Semantic Stability

- Models require stable meaning. If "batch start," "mixing phase," "hold time," or "CPP excursion" are defined differently across systems or sites, model training becomes a function of representational artifacts rather than manufacturing reality.
- The ontology construct introduced in Section 5 addresses this by externalizing meaning as a governed contract, enabling semantic interoperability across discovery, development, and manufacturing.

10.6.2. Event Coherence and Temporal Integrity

- Many manufacturing AI use cases are fundamentally temporal: they depend on the ordering of events, state transitions, and time-in-state. Without event coherence, models cannot reliably distinguish causal signals from correlated noise.
- Event-centric execution representation (Section 6) provides this temporal backbone, allowing inference to be anchored in execution reality rather than inferred timestamps.

10.6.3. Context Binding and Feature Reliability

- In manufacturing, a measurement is rarely meaningful without context. A temperature reading, pressure signal, or assay result becomes relevant only when tied to the correct step, material state, equipment role, and control strategy.
- Binding context at generation reduces feature instability and improves model portability across lifecycle phases and sites.

10.6.4. Provenance, Lineage, and Transformation Transparency

- AI systems are sensitive to data transformations, augmentations, and label construction. If these transformations are not traceable, model training becomes irreproducible, and inference becomes difficult to defend.
- The NIST AI RMF (NIST, 2023) and ISO/IEC 42001 (ISO/IEC, 2023) both emphasize the importance of data provenance and transformation integrity for trustworthy AI, while GAMP 5 (ISPE, 2022) provides the GxP-specific framework for validating the computerized systems that manage these transformations.

10.6.5. Evidence Models for GxP Credibility

- In regulated contexts, models must be validated and maintained. ISPE's AI in pharmaceutical manufacturing guidance (ISPE, 2025) emphasizes that assurance of integrity and quality of outputs is essential and must be considered early to avoid pilot stagnation. Rudin (2019) argues more broadly that high-stakes domains should prioritize

inherently interpretable models over post-hoc explanations of black boxes—a position with direct implications for GxP contexts where decision rationale must be auditable.

- This paper's evidence layer concept supports that requirement by structuring data and context so it remains suitable for review and compliance consumption.

10.7. AI Enablement Across Key Manufacturing Domains

- To illustrate how these prerequisites manifest, AI enablement can be grouped into several archetypal domains. The purpose here is not to enumerate use cases exhaustively, but to show how architecture determines feasibility.

10.7.1. PROCESS MONITORING AND ANOMALY DETECTION

- These applications benefit from high-resolution signals locally but require event and context coherence to scale beyond single-unit deployments. Without contextual binding, anomaly outputs are difficult to interpret and operationalize.

10.7.2. PREDICTIVE QUALITY AND DEVIATION RISK FORECASTING

- These models require strong semantic stability and evidence traceability because they inform quality decisions and may be subject to heightened scrutiny. Inconsistent definitions of quality attributes or process states undermine credibility.

10.7.3. TECH TRANSFER LEARNING AND CROSS-SITE GENERALIZATION

- Cross-site learning is a direct beneficiary of the IDDC. If execution events and context are comparable, models can generalize and support transfer readiness. If not, each site becomes an isolated training domain and AI value cannot compound.

- This mirrors Section 9's argument that tech transfer is a stress test of architecture; it is also a stress test of AI generalizability.

10.7.4. KNOWLEDGE RETRIEVAL AND DECISION SUPPORT

- Generative AI, retrieval-augmented generation, and agentic patterns are increasingly discussed across the industry, including within BioPhorum programming.

- However, these approaches are only as reliable as the semantic organization of the underlying knowledge. Without controlled meaning, provenance, and evidence structures, generative outputs become difficult to trust in regulated decisions.

- The BioPhorum data characteristics framework (BioPhorum, 2026) reinforces this point: knowledge retrieval and decision support applications require the highest levels of semantic connection and contextualization — precisely the machine understandability characteristics that depend on ontological grounding. Without the Semantic Layer that ontology provides, generative and agentic AI outputs in regulated contexts cannot achieve the interpretability and traceability that quality organizations require for trust.

10.8. Failure Modes When AI Is Pursued Without Architectural Readiness

- A key risk in life sciences digital strategies is the inversion of sequencing: pursuing AI before establishing the foundations that make it trustworthy and scalable. When this occurs, predictable failure modes follow:

- Models become dependent on local feature engineering and informal context knowledge
- Outputs are not reproducible across time, sites, or lifecycle phase

- Validation becomes difficult because inputs and transformations cannot be defended
- Quality organizations distrust results due to weak evidence lineage
- Programs remain trapped in pilots, reinforcing skepticism about AI value

- These are not primarily model failures. They are architectural failures.

10.9. Implications for Program Design and Governance

- AI enablement requires a governance model that matches its lifecycle nature. Regulatory and standards bodies increasingly emphasize that AI must be governed through continuous risk management, credibility assessment, and lifecycle controls. The NIST AI RMF (NIST, 2023), ISO/IEC 42001 (ISO/IEC, 2023), and the EU AI Act (European Union, 2024) converge on a common expectation: AI governance must be risk-proportionate, lifecycle-aware, and embedded in organizational management systems rather than treated as a one-time validation exercise.

- Within the framing of this paper, governance should focus on:

- defining and approving contexts of use
- ensuring data readiness is measurable and observable, not assumed
- treating semantic constructs and evidence models as controlled lifecycle assets
- embedding accountability for model operation, monitoring, and change management
- maintaining transparency on drift, retraining triggers, and performance boundaries

- This aligns with ISPE's governance and QA framing for AI in GxP contexts (ISPE, 2025), which emphasizes sustained assurance rather than one-time validation events. This paper will provide a more detailed discussion of Governance in Section 12.

10.10. Future Directions: From Decision Support to Agentic Manufacturing Intelligence

- The preceding subsections establish that AI enablement in regulated manufacturing is fundamentally an architectural capability, not a technology deployment. The BioPhorum data characteristics framework validates that AI-ready data requires semantic consistency, contextual binding, and traceable provenance — capabilities that the IDDC provides by design. The failure modes described above demonstrate what happens when AI is pursued without these foundations. This subsection looks forward to a frontier that current architectural investments make possible: the emergence of agentic AI systems that can perceive, reason about, and act within manufacturing environments with increasing autonomy — not as replacements for human judgment, but as participants in a collaborative intelligence that exceeds what either humans or machines achieve alone.

- Without an architectural foundation that preserves manufacturing intent, context, and evidence as navigable, queryable structures, every deployment of an AI agent in manufacturing requires bespoke engineering. Data products must be constructed manually for each use case. Integration interfaces must be purpose-built to expose the specific data each agent needs. Explicit instructions must encode what matters, what relates to what, and why — effectively recreating by hand, for each agent, the contextual knowledge that a well-constructed IDDC would provide as architectural infrastructure. This is the same lifecycle discontinuity described in Section 4, now manifesting at the AI layer: the fragmentation of meaning that impedes technology transfer and batch disposition equally impedes the scaling of intelligent systems. The result is what industry practitioners increasingly recognize as pilot purgatory — so much effort devoted to constructing the prerequisites for

each AI application that organizations exhaust their capacity for experimentation before reaching operational value.

- With a functioning IDDC, the architectural situation inverts. The four layers — Semantic, Event, Context, and Evidence — collectively function as a navigable knowledge fabric that AI agents can traverse. An agent operating within this fabric can discover relationships between process parameters and quality attributes by querying the ontology. It can understand the temporal structure of manufacturing execution through the event layer. It can access the contextual bindings that explain why a measurement matters for a specific batch, material, or patient. It can evaluate the evidentiary record that supports or challenges a disposition decision. Critically, the ontology itself becomes an active enabler — not merely a reference that humans consult, but a machine-readable map that agents use to navigate systems and data, understand what they are observing, and determine appropriate actions within governed boundaries. The IDDC transforms agent deployment from a bespoke engineering exercise into an architectural capability: each new agent benefits from the same governed semantic foundation, the same contextual infrastructure, and the same evidentiary structures — compounding value rather than fragmenting it.
- This vision is grounded in Malone's (2018) framework of collective intelligence and superminds — systems in which humans and machines collaborate to achieve outcomes that neither can achieve independently. In manufacturing, this collaboration is not hypothetical; it is already practiced in embryonic form. The AMP case study (Appendix A) illustrates the pattern concretely: human operators and a digital platform collaborate to disposition alarms, with the platform providing unified presentation and contextual evidence while humans exercise judgment on significance and response. With architectural maturity, this collaboration deepens. AI models could begin to pre-disposition certain alarms within defined boundaries — not replacing human judgment, but filtering noise, surfacing patterns, and presenting evidence in forms that accelerate human decision-making. This is the first step on the path from decision support to agentic operation, grounded in something that has already been built and is generating operational evidence.
- The IDDC can be understood as a fabric from which different threads can be pulled — threads connecting systems, people, processes, and machines — and woven together to identify and surface opportunities or risks within operations and data. Organizations that invest in this fabric build architectural readiness: the capacity to adopt proven AI approaches rapidly as they emerge, because the data foundations, semantic structures, and evidentiary infrastructure are already in place. Without this readiness, the gap between observing what AI can achieve elsewhere and executing it internally becomes insurmountable. The fast follower strategy — deliberately adopting proven innovations rather than pioneering them — degrades from a strategic choice into involuntary stagnation, because each new AI initiative must first reconstruct the architectural prerequisites that a well-designed IDDC would have provided from the outset.
- Early industry examples offer partial validation of this trajectory. Organizations that have publicly described successful AI scaling in pharmaceutical manufacturing consistently emphasize that architectural foundations preceded analytical deployment — investment in unified data access, semantic consistency, and governed context created the conditions under which AI models could generalize beyond individual sites or use cases. These examples demonstrate a pattern consistent with the IDDC argument: AI value compounds when it operates on architecturally coherent data, and fragments when each application must construct its own contextual foundation.

- It should be stated clearly that agentic AI in regulated manufacturing remains a future direction, not a present reality. No published example demonstrates autonomous AI agents operating within GMP manufacturing environments with the degree of contextual understanding and governed authority described here. The contribution of this subsection is architectural logic, not empirical validation: if the IDDC provides what the preceding sections argue it provides — preserved intent, coherent events, bound context, and structured evidence — then the architectural prerequisites for agentic capability are satisfied by design. The BioPhorum framework validates that these are the prerequisites industry consensus identifies. What remains is the engineering, governance, and organizational maturation required to realize this potential — a trajectory that current investments in platforms like AMP, ontology infrastructure, and data architecture are actively advancing.

Key Takeaway

AI enablement in life sciences manufacturing is not achieved through technology adoption alone. It is achieved when the architecture preserves meaning, event coherence, contextual binding, and quality evidence across the lifecycle. The IDDC provides the architectural foundation that makes AI credible and scalable, transforming AI from isolated experimentation into a governed, lifecycle-capable decision-support capability suitable for regulated environments.

-

11. REFERENCE ARCHITECTURE FOR LIFECYCLE-ALIGNED DIGITAL EXECUTION

- The preceding sections have established that sustainable digital value in life sciences manufacturing depends on preserving manufacturing intent, execution context, and quality evidence across the full product lifecycle. This section synthesizes those concepts into a reference architectural model, following established reference architecture practice as described in enterprise architecture frameworks such as TOGAF (The Open Group, 2018) and Zachman's (1987) foundational work on information systems architecture, that illustrates how lifecycle continuity can be achieved without prescribing specific technologies, platforms, or organizational structures.
- The purpose of this reference architecture is not to define a target system landscape—an approach that would contradict the technology-agnostic principles of established EA frameworks (The Open Group, 2018)—but to articulate the architectural responsibilities required to support scalable execution, digital tech transfer, continuous batch disposition, credible AI enablement, and more.

11.1. Architectural Design Principles

- The reference architecture presented here is governed by several core design principles derived directly from the analysis in earlier sections:
- Lifecycle first
- Architectural constructs must span discovery, development, manufacturing, and transfer, rather than optimizing for a single phase.

- Intent before execution
- Manufacturing intent must be explicitly defined and governed independently of any execution environment.
- Meaning over movement
- Preserving semantic meaning and context is more critical than maximizing data flow or centralization.
- Evidence by design
- Quality-relevant evidence must be generated continuously as execution occurs, not reconstructed retrospectively.
- Decoupled but coherent
- Systems may evolve independently, but architectural meaning must remain stable.

- These principles ensure that the architecture remains durable as products, sites, and technologies change. They reflect the separation of concerns that is central to enterprise architecture practice (Zachman, 1987; The Open Group, 2018) while grounding each principle in the specific regulatory and operational requirements of life sciences manufacturing. Figure 7 provides a visual representation of this reference architecture, illustrating how the architectural domains relate to one another and to the IDDC.

-

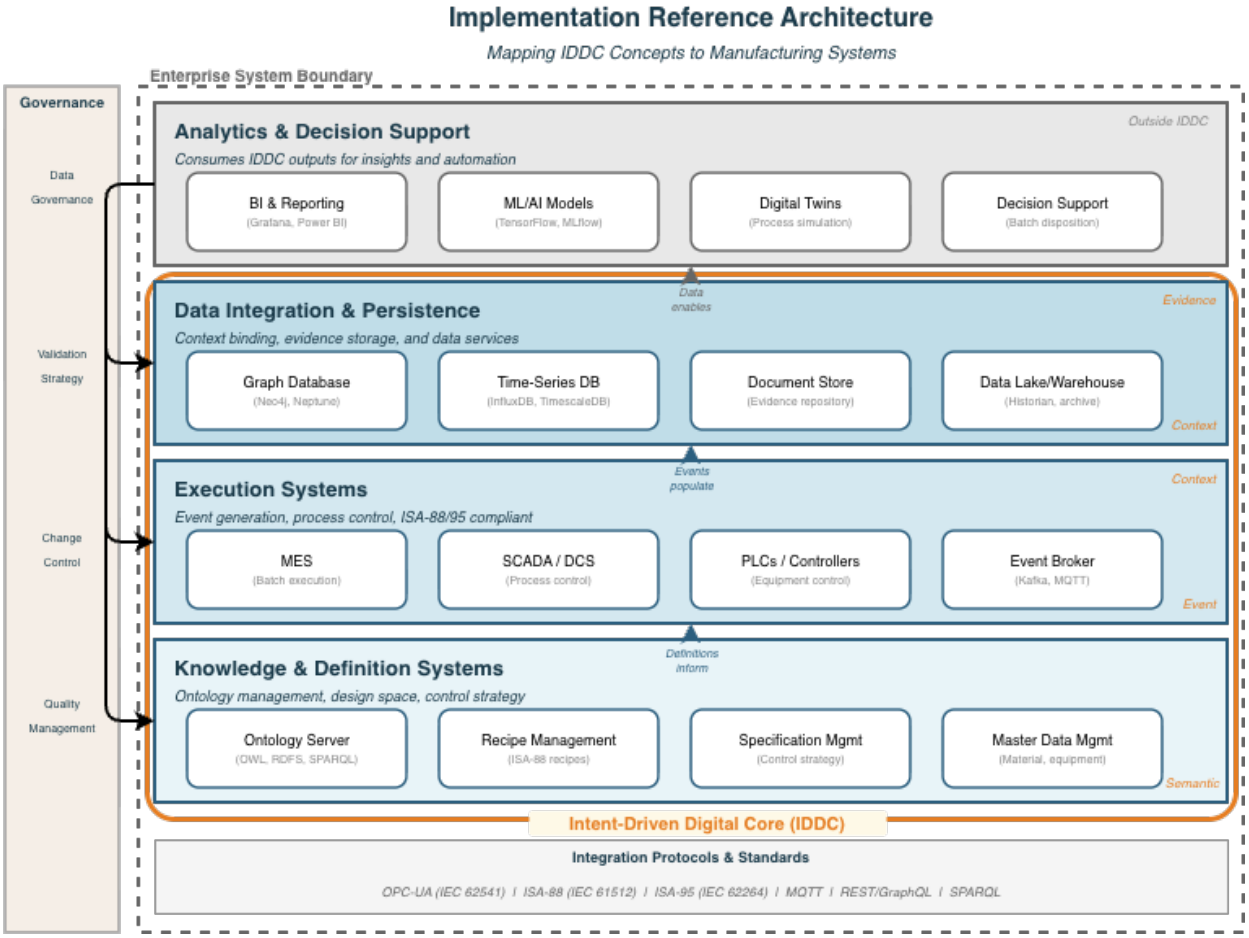

- Figure 7. Implementation Reference Architecture mapping enterprise systems to IDDC concepts. The three lower layers implement IDDC capabilities; Analytics consumes IDDC outputs. See Figure 2 for conceptual IDDC architecture.

- As shown in Figure 7, the reference architecture positions the IDDC as the central substrate connecting the Manufacturing Intent Domain, Execution Domain, Context and Evidence Domain, Integration Domain, and Governance Domain. The layered model demonstrates how intent flows from definition through execution to evidence, with each domain maintaining clear responsibilities while contributing to lifecycle continuity. Importantly, the figure illustrates that the architecture is technology-agnostic: specific tools and platforms can be selected based on organizational context, provided they fulfill the architectural responsibilities defined for each domain. This approach enables incremental adoption and respects existing system investments while establishing the structural foundation for scalable lifecycle capabilities.

11.2. Core Architectural Domains

- At a high level, the reference architecture can be understood as a set of interrelated domains, each with a distinct responsibility.

11.2.1. MANUFACTURING INTENT DOMAIN

- This domain is responsible for defining and governing manufacturing intent across lifecycle phases, operationalizing the concepts of design space and control strategy from ICH Q8(R2) (ICH, 2009) and the lifecycle knowledge management requirements of ICH Q10 (ICH, 2008). It includes explicit representations of:

- process objectives and constraints,
- critical quality and process attributes,
- control strategies and acceptance criteria,
- relationships between materials, equipment roles, and process stages.

- Importantly, intent is versioned and traceable, enabling controlled evolution as knowledge matures from development through commercialization. To the extent possible, these representations should be structured data. Where documents are needed for consumption by humans, they should be generated deterministically for consistency.

11.2.2. EXECUTION DOMAIN

- The execution domain encompasses control systems, manufacturing execution systems, laboratory systems, and other platforms responsible for carrying out manufacturing activities. These systems are typically organized according to ISA-88 (IEC 61512) procedural control hierarchies and ISA-95 (IEC 62264) activity models, which provide the structural framework within which execution events are generated and interpreted.

- Within the reference architecture, execution systems are not required to be uniform or centrally governed. Their architectural responsibility is to:

- expose meaningful execution events,
- associate those events with the applicable intent in effect at the time,
- preserve temporal coherence and state transitions.

- This allows execution reality to be captured faithfully without constraining local implementation choices.

11.2.3. CONTEXT AND EVIDENCE DOMAIN

- The context and evidence domain binds execution events to the semantic and operational context necessary for interpretation. This includes:

- process and phase context,
- material and batch lineage,
- equipment configuration and role context,
- quality expectations and applicable limits.

- Evidence generated within this domain is structured to support downstream quality review, regulatory inspection, and analytical consumption without additional reconciliation.

11.2.4. INTEGRATION AND INTEROPERABILITY DOMAIN

- Rather than focusing on point-to-point integrations, this domain ensures that systems interoperate against shared meaning. Ontologies, canonical event structures, and governed interfaces enable systems to exchange information without semantic loss. This approach reflects Evans's (2003) domain-driven design principle of a "ubiquitous language"—a shared vocabulary that bridges technical and domain experts—and applies Hohpe and Woolf's (2003) enterprise integration patterns to ensure that messaging and data exchange preserve semantic fidelity across system boundaries.

- This domain is where lifecycle portability is achieved, supporting both technology transfer and AI generalization.

11.2.5. GOVERNANCE AND STEWARDSHIP DOMAIN

- The reference architecture assumes that architectural continuity is actively stewarded. Governance mechanisms define:

- ownership of intent and semantic definitions,
- accountability for data products and evidence quality,
- change management for evolving processes and models,
- lifecycle oversight for analytics and AI capabilities.

- Governance is positioned as an enabling function, ensuring consistency without constraining execution.

11.3. THE ROLE OF THE INTENT-DRIVEN DIGITAL CORE

- The IDDC sits at the center of this reference architecture as the unifying construct that binds intent, execution, context, and evidence across domains.

- Rather than acting as a monolithic system, the IDDC represents a persistent architectural foundation that:

- maintains semantic continuity as processes evolve,
- preserves event coherence across systems and sites,
- accumulates quality-relevant evidence progressively,
- enables downstream capabilities to consume trustworthy representations of manufacturing reality.

- In this model, analytics, AI, digital twins, batch disposition, and tech transfer are consumers of the core, not extensions of it. This distinction ensures that value compounds as capabilities are added, rather than fragmenting with each new initiative. In enterprise architecture terms, the IDDC functions as the canonical data model (Hohpe & Woolf, 2003) that decouples consumers from the operational systems that generate data, enabling new capabilities to be added without redesigning the foundational infrastructure.

11.4. Architectural Outcomes Across the Lifecycle

- When implemented coherently, the reference architecture supports several critical outcomes:
- Scalable technology transfer, enabled by portable intent and comparable evidence
- Continuous batch disposition, supported by progressive evidence accumulation
- Credible analytics and AI, grounded in semantically stable, contextualized data
- Reduced manual reconciliation, shifting effort toward informed judgment
- Improved regulatory confidence, driven by traceability and evidence integrity
- These outcomes emerge from architectural alignment, not from the optimization of individual systems.

11.5. From Reference Model to Organizational Adoption

- The reference architecture is intentionally abstract. Organizations will instantiate it differently based on their operating models, technology landscapes, and regulatory strategies. The ISPE Pharma 4.0 Digital Maturity Model (ISPE, 2023) provides a useful assessment framework for thisinstantiation, enabling organizations to evaluate their current capabilities against the architectural responsibilities defined here. However, successful adoption consistently requires:
- explicit recognition of intent and evidence as lifecycle assets,
- alignment of people and roles with architectural responsibilities,
- governance models that prioritize continuity over local optimization,
- phased implementation that strengthens foundations before advanced analytics and AI.
- Importantly, organizations need not reach full maturity across all domains simultaneously. Value can be realized incrementally, provided architectural coherence is preserved.

> ***Key Takeaway***
>
> A lifecycle-aligned digital architecture is defined by responsibilities, not tools. By centering intent, preserving context, and generating evidence by design, the reference architecture described here provides a durable foundation for scalable manufacturing, digital tech transfer, continuous quality oversight, and trustworthy AI across the life sciences manufacturing lifecycle.

-

12. GOVERNANCE AND OPERATING MODEL IMPLICATIONS

- The architectural strategy described in this paper cannot be sustained through technology decisions alone. Preserving manufacturing intent, contextual integrity, and evidence across the lifecycle requires an operating model that aligns accountability, decision rights, and stewardship with the architecture itself. Governance, in this context, is not a compliance overlay or a project control mechanism; it is how architectural continuity is maintained as products, processes, and capabilities evolve. While established data governance frameworks such as DAMA-DMBOK (DAMA International, 2017) provide foundational principles for data management, the governance demands of an intent-driven architecture

extend beyond data stewardship to encompass semantic governance, evidence lifecycle management, and cross-functional accountability for manufacturing intent.

- This section examines the governance and operating model implications of an intent-driven digital architecture, with particular attention to how emerging AI capabilities fundamentally change what must be governed and how.

12.1. Lifecycle Ownership of Manufacturing Intent

- In traditional operating models, ownership of manufacturing intent is fragmented. Development teams define process understanding, manufacturing teams focus on execution, and quality organizations interpret evidence retrospectively. As products move across lifecycle phases, intent is re-expressed through documents, specifications, and system configurations, often without a single accountable owner. ICH Q10 (ICH, 2008) recognizes this fragmentation as a barrier to effective pharmaceutical quality systems, while ICH Q12 (ICH, 2019) attempts to address it through the concept of established conditions that create explicit lifecycle ownership of critical process parameters.
- An intent-driven architecture requires a different model: manufacturing intent must be owned as a lifecycle asset. Governance mechanisms must ensure that intent is explicitly defined, versioned, and stewarded as it evolves from discovery through commercialization and technology transfer. This does not imply centralization of authority, but rather clarity of accountability for how intent is represented and preserved. Mintzberg's (1979) analysis of organizational design demonstrates that effective coordination can be achieved through standardization of outputs and skills rather than centralization of decision-making—a principle directly applicable to governing manufacturing intent across distributed organizations.
- Lifecycle ownership of intent enables several critical outcomes:
- controlled evolution of process definitions as knowledge matures,
- consistent interpretation of expectations across sites and partners,
- traceability between development assumptions and manufacturing execution,
- and alignment between quality decisions and original process objectives.
- Without explicit ownership, intent inevitably degrades into locally interpreted artifacts, undermining the architectural foundation.
- Cross-Functional Accountability by Design
- The IDDC collapses traditional boundaries between IT, OT, and Quality by design. Semantic consistency, event coherence, and evidence integrity cannot be maintained within functional silos. As a result, governance models that rely on functional handoffs or system-centric accountability become liabilities.
- Effective governance in this model is cross-functional by default. Accountability is aligned to architectural responsibilities rather than organizational units. For example:
- intent stewardship spans development, manufacturing, and quality,
- data product ownership crosses operational and analytical domains,
- evidence integrity involves both execution and quality functions.
- This shift does not eliminate functional expertise. Instead, it requires governance structures that explicitly manage interdependencies and shared ownership. Decision rights must be

clear not only within functions, but across them, particularly where architectural consistency is at stake.

12.2. The Evolving Role of Centers of Excellence

- Centers of Excellence (COEs) are commonly established to accelerate digital initiatives, standardize tools, or concentrate expertise. In many organizations, however, COEs become control points that inadvertently slow delivery or fragment ownership by pulling responsibility away from operational teams. Galbraith (2014) identifies this as a common failure mode of lateral organizational mechanisms: structures intended to enable coordination instead become bottlenecks when they accumulate decision rights without corresponding accountability for outcomes.
- Within an intent-driven architecture, the role of the COE must evolve. COEs are most effective when positioned as enabling functions, responsible for:
- defining and maintaining architectural standards,
- stewarding shared semantic models and patterns,
- enabling reuse of capabilities across the organization,
- supporting education and architectural literacy.
- They should not function as solution factories or approval bottlenecks. The objective is architectural coherence, not centralized execution. When COEs are aligned to capability stewardship rather than project delivery, they reinforce continuity rather than constrain it. The ISPE Pharma 4.0 Operating Model (ISPE, 2019) advocates precisely this positioning, recommending that Centers of Excellence function as enabling capabilities that develop and maintain organizational standards while empowering operational teams to execute within governed boundaries.
- From Project Delivery to Capability Stewardship
- Many digital programs are governed as projects with defined start and end points. While effective for system implementation, this model is poorly suited to architectural constructs that must persist and evolve over time.
- An intent-driven digital architecture demands a shift from project delivery to capability stewardship. Core architectural elements such as semantic definitions, event models, evidence structures, and AI models are not “completed” at project close. They require ongoing stewardship, change management, and fitness-for-purpose assessment. GAMP 5 (ISPE, 2022) provides the GxP-specific framework for this lifecycle approach to computerized system governance, while Kotter (1996) emphasizes that sustainable change requires institutional anchoring—embedding new practices in organizational culture and operational processes rather than relying on project-level momentum.
- Operating models must therefore:
- recognize architectural capabilities as long-lived assets,
- fund stewardship alongside delivery,
- and define success in terms of sustained value rather than milestone completion.
- This shift is particularly important for AI-enabled capabilities, which are inherently dynamic and sensitive to drift in data, process, and context.

12.3. Governance Implications of AI Enablement

- AI fundamentally changes the governance landscape. Unlike traditional analytics, AI models learn from data, adapt over time, and can influence decisions at scale. As a result, governance must extend beyond data access and model validation to include lifecycle assurance of meaning, context, and risk. The convergence of the NIST AI RMF (NIST, 2023), ISO/IEC 42001 (ISO/IEC, 2023), and the EU AI Act (European Union, 2024) establishes an emerging governance framework that requires risk-proportionate oversight, continuous monitoring, and explicit accountability for AI system performance throughout deployment.
- Key governance implications include:
- explicit definition and approval of AI contexts of use,
- accountability for training data provenance and semantic fitness,
- monitoring for concept drift and process change,
- and integration of AI outputs into quality and operational decision frameworks.
- Critically, AI governance cannot be bolted on after models are built. It must be embedded in the operating model that governs the IDDC. When intent, events, context, and evidence are preserved architecturally, AI governance becomes tractable. When they are not, AI risk escalates rapidly and confidence erodes. ICH Q9(R1) (ICH, 2023) reinforces this point: quality risk management must be systematic, science-based, and proportionate to risk—principles that apply with particular force to AI-enabled capabilities where the scope and pace of decision-making exceed what traditional governance models were designed to oversee.

12.4. Governance Without Over-Specification

- A central risk in governance design is over-specification. Excessive control structures can stifle innovation and slow adoption, particularly in fast-moving development environments. The governance model implied by this architecture is therefore principle-driven rather than prescriptive.
- Effective governance focuses on:
- what must remain consistent (intent, meaning, evidence),
- where flexibility is acceptable (local execution, tooling),
- and how change is managed transparently.
- This balance allows organizations to scale digital capabilities, including AI, without sacrificing architectural integrity or regulatory confidence.

Key Takeaway

Architecture only delivers sustained value when it is matched by an operating model that governs intent, meaning, and evidence across the lifecycle. In an era of AI-enabled manufacturing, governance must evolve from project control and system oversight to capability stewardship and lifecycle accountability. Organizations that align governance to architectural continuity position themselves to compound digital value rather than repeatedly re-solve the same problems.

-

13. CONCLUSION: FROM DIGITAL PROJECTS TO DIGITAL CONTINUITY

- The life sciences industry has invested heavily in digital technologies across discovery, development, and manufacturing. Yet the uneven realization of value from these investments suggests that the core challenge is not ambition, tooling, or effort, but architecture. Digital initiatives that are conceived and delivered as isolated projects struggle to scale, integrate, and endure because they do not preserve manufacturing intent, context, and evidence as products and processes evolve.
- This paper has argued that digital continuity, rather than digital sophistication, is the defining requirement for sustainable value. Continuity is achieved not through system consolidation or advanced analytics alone, but through an architectural strategy that explicitly governs how meaning is defined, how execution reality is captured, and how evidence is accumulated across the lifecycle.

13.1. Why Architecture Must Precede Analytics and AI

- Advanced analytics and artificial intelligence are increasingly central to the future of pharmaceutical development and manufacturing. However, when deployed atop fragmented data landscapes, they amplify inconsistency rather than insight. Models trained on reconstructed, weakly contextualized data inherit the ambiguity and bias of their inputs —the "hidden technical debt" that Sculley et al. (2015) identify as the systemic risk of deploying machine learning on poorly governed data foundations—undermining trust and limiting applicability beyond narrow use cases.
- By contrast, an intent-driven digital architecture establishes the conditions under which analytics and AI can compound in value. When manufacturing intent is explicit, execution events are coherent, context is bound at the point of generation, and evidence is preserved by design, analytical and AI capabilities consume a stable representation of reality. In this sequence, architecture is not a prerequisite that delays innovation, but the enabler that allows innovation to scale responsibly.
- The BioPhorum consortium's data characteristics framework reinforces this sequence empirically. Its four groups of AI-ready data requirements — data observability, data quality, data availability, and machine understandability — each presuppose capabilities that the IDDC's architectural layers provide: observability depends on coherent event capture (Event layer), quality depends on governed semantic definitions (Semantic layer), availability depends on contextualized data products (Context layer), and machine understandability depends on traceable, structured evidence (Evidence layer). Without this architectural foundation, organizations pursue AI readiness as a data remediation exercise rather than an architectural outcome.

13.2. Continuity as the Mechanism for Compounding Value

- One of the most persistent frustrations in digital transformation is the need to repeatedly re-solve the same problems at each site, for each product, or at every lifecycle transition. This repetition is not inevitable. It is a consequence of architectures that do not carry meaning forward.
- Digital continuity changes this dynamic. When intent, context, and evidence persist across phases and sites, knowledge compounds. Technology transfer becomes repeatable rather than bespoke. Batch disposition shifts from episodic reconstruction to continuous evaluation. AI models generalize across processes and facilities rather than being retrained

in isolation. The BioPhorum data characteristics framework (BioPhorum, 2026) validates this at the consortium level: AI-ready data requires semantic consistency, contextual binding, and traceable provenance—characteristics that compound naturally within an intent-driven architecture.

- In this model, value accrues not from individual digital projects, but from the accumulation and reuse of architectural assets.
- The alarm management platform described in Appendix A illustrates this compounding pattern concretely. What began as event capture — unifying fragmented alarm data into a single, coherent record — enabled contextualized alarm disposition by operators and quality reviewers, which in turn generates progressive evidence for batch release decisions. Each architectural layer creates the conditions for the next. With sufficient maturity, the same platform positions the organization to deploy AI models that pre-disposition certain alarms within defined boundaries — not as a standalone analytics initiative, but as a natural extension of the architectural investment already made.

13.3. What Organizations Gain from Intent-Driven Design

- Organizations that adopt an intent-driven digital architecture gain more than operational efficiency. They gain:
- confidence in decision-making, grounded in traceable evidence,
- agility in scaling processes and transferring technology,
- credibility in deploying AI within regulated environments,
- and resilience as systems, products, and regulatory expectations evolve.
- Equally important, they shift the focus of digital transformation away from technology selection and toward architectural stewardship. This reframing aligns people, processes, and systems around a shared understanding of manufacturing reality—what Malone (2018) characterizes as the foundation for collective intelligence in organizations where humans and machines collaborate to produce outcomes neither achieves independently.

13.4. Limitations

- **Scope Boundaries.** This paper addresses digital architecture strategy for Chemistry, Manufacturing, and Controls (CMC) operations in pharmaceutical and biopharmaceutical manufacturing. The framework does not address discovery research, clinical trial operations, or commercial distribution and supply chain. While the architectural principles articulated here — preserving intent, binding context at execution, accumulating evidence progressively — may apply to adjacent domains, the specific regulatory grounding (ICH Q-series, GAMP 5, FDA and EMA guidance), industry examples, and consortium validation are drawn exclusively from manufacturing contexts. Applicability beyond this scope should not be assumed without independent evaluation.
- **Methodological Boundaries.** The Intent-Driven Digital Core is developed through design science methodology — constructing and evaluating an architectural artifact against requirements derived from industry practice, regulatory expectations, and consortium consensus. It has not been subjected to controlled empirical validation across multiple manufacturing sites or organizations. The empirical evidence presented — the alarm management platform (Appendix A), the ontology deployment architecture (Appendix B), and the technology transfer analysis (Section 9) — demonstrates architectural feasibility and internal coherence rather than statistical generalizability. Each evidence stream

illustrates how IDDC principles operate in practice; none constitutes a formal comparative study.

- **Generalizability.** The regulatory and standards landscape that shapes this framework — GMP requirements, ALCOA+ data integrity expectations, ICH quality guidelines — is specific to pharmaceutical and biopharmaceutical manufacturing. Other regulated industries (medical devices, food safety, aerospace) share analogous requirements for traceability, data integrity, and lifecycle governance, but operate under different regulatory regimes with different enforcement mechanisms. Whether the IDDC's specific architectural patterns transfer to these contexts, or whether the framework requires substantive adaptation, remains an open question.
- **Economic Analysis.** This paper does not include cost-benefit analysis, return-on-investment modeling, or economic justification for IDDC adoption. The argument is architectural: that lifecycle discontinuity is a structural problem requiring an architectural response. Organizations evaluating adoption will need to develop their own economic models, accounting for implementation scope, existing system landscape, organizational readiness, and the opportunity cost of continued fragmentation.

13.5. Future Research Directions

- **Empirical Validation.** The most immediate research need is multi-site empirical validation of IDDC architectural principles. A cross-company study — conducted within a consortium context such as BioPhorum, the Pistoia Alliance, or NIIMBL — could evaluate whether organizations that adopt intent-driven architectural patterns achieve measurable improvements in technology transfer efficiency, batch disposition cycle time, and AI deployment success rates compared to organizations relying on conventional integration approaches. Such a study would require standardized maturity assessment instruments and longitudinal observation across implementation phases.
- **Architectural Maturity Metrics.** This paper argues qualitatively that architectural coherence compounds value over time. Developing quantitative metrics for IDDC maturity — analogous to capability maturity models but specific to manufacturing digital architecture — would enable organizations to benchmark progress and prioritize investment. Metrics should address each IDDC layer independently (semantic completeness, event coherence, context binding coverage, evidence accumulation rate) and their integration collectively, distinguishing between organizations that have implemented individual layers and those that have achieved cross-layer architectural coherence.
- **Cross-Industry Applicability.** Testing IDDC principles in adjacent regulated industries — medical devices under FDA 21 CFR Part 820, food manufacturing under FSMA, aerospace under AS9100 — would determine whether the framework's regulatory grounding is industry-specific or whether the core architectural argument generalizes. The hypothesis is that the need to preserve design intent across lifecycle phases is universal in regulated manufacturing; the specific mechanisms and governance structures may require domain adaptation.
- **AI Governance in Architecturally Coherent Environments.** Section 10 and the emerging vision for agentic manufacturing intelligence raise governance questions that this paper identifies but does not fully resolve. Research is needed on governance frameworks designed specifically for AI operating within intent-driven architectures — where models have access to contextualized, semantically grounded data rather than isolated datasets. The governance challenge shifts from data quality assurance to intent preservation

assurance: ensuring that AI-driven decisions remain traceable to the manufacturing intent they are meant to serve.

13.6. A Call for Architectural Discipline

- The transition from digital projects to digital continuity requires discipline. It requires resisting the temptation to optimize locally at the expense of lifecycle coherence. It requires investing in foundational constructs that may be less visible than dashboards or models, but far more consequential over time.
- The IDDC described in this paper is not a prescription for a specific platform or operating model. It is a way of thinking about architecture as the mechanism by which intent becomes execution, execution becomes evidence, and evidence becomes knowledge.
- As life sciences manufacturing continues to evolve in complexity and scale, the organizations that succeed will be those that treat architecture not as a technical artifact, but as a strategic asset.
-

APPENDIX: GLOSSARY OF SUPPORTING TERMS

Term	Definition
Batch Disposition	The formal quality decision to release or reject a manufactured batch based on review of execution records, test results, deviations, and supporting evidence.
Control Strategy	A planned set of controls derived from process understanding that ensures process performance and product quality across the manufacturing lifecycle.
Critical Quality Attribute (CQA)	A physical, chemical, biological, or microbiological property that must be maintained within defined limits to ensure product quality.
Critical Process Parameter (CPP)	A process parameter whose variability has a direct and significant impact on a critical quality attribute and must be controlled.
Electronic Batch Record (EBR)	A digital representation of manufacturing execution and quality data used to support batch review and disposition.
GxP	A collection of good practice guidelines that ensure products are safe, effective, and compliant with regulatory requirements.
ALCOA+	A set of principles ensuring data integrity by requiring data to be attributable, legible, contemporaneous, original, accurate, complete, consistent, enduring, and available.
Data Integrity	The completeness, consistency, and accuracy of data throughout its lifecycle, including its creation, modification, maintenance, and archival.
Data Provenance	Documentation of the origin, lineage, and transformations applied to data throughout its lifecycle.
Event-Driven Architecture	A design pattern in which system behavior and data capture are structured around discrete events and state changes rather than static records.
Interoperability	The ability of systems to exchange information and interpret it consistently without loss of meaning.
Ontology	A formal, machine-readable representation of concepts, relationships, and definitions within a domain, used to enable semantic consistency.
Semantic Interoperability	The ability of systems to exchange data with shared understanding of meaning, not merely shared syntax.

Term	Definition
Unified Namespace (UNS)	A logical data architecture pattern that exposes operational data in a consistent, real-time structure, commonly used in manufacturing contexts.
Technology Transfer	The process of transferring manufacturing knowledge, processes, and controls from development to manufacturing or between manufacturing sites.
Lifecycle Phase	A distinct stage in the product lifecycle, such as discovery, development, manufacturing, or technology transfer.
Operating Model	The combination of structures, processes, roles, and governance mechanisms by which an organization delivers value.
Center of Excellence (COE)	A centralized group established to provide expertise, standards, and enablement for specific capabilities or disciplines.
Change Management	The structured approach to managing modifications to processes, systems, or controls to maintain compliance and performance.
Advanced Analytics	Analytical techniques beyond descriptive reporting, including statistical modeling, optimization, and predictive methods.
Artificial Intelligence (AI)	Computational techniques that enable systems to perform tasks requiring human-like reasoning, learning, or pattern recognition.
Machine Learning (ML)	A subset of artificial intelligence in which models learn patterns from data rather than being explicitly programmed.
Model Validation	The documented process of ensuring that a model performs as intended within its defined context of use.
Context of Use (AI)	The specific conditions, assumptions, and decision boundaries under which an AI model is intended to operate.
Explainability	The degree to which the internal logic and outputs of a model can be understood by humans.
Governance	The framework of decision rights, accountability, and oversight used to ensure consistency, compliance, and strategic alignment.
Risk-Based Approach	A method of prioritizing controls and oversight based on the potential impact to quality, safety, or compliance.
Regulatory Compliance	Adherence to applicable laws, regulations, and guidance governing product development and manufacturing.
Inspection Readiness	The state of preparedness to demonstrate compliance and data integrity during regulatory inspection.

- # APPENDIX: BIBLIOGRAPHY
- **Regulations, regulators, and GMP data integrity**

- European Commission. (2011). EudraLex, Volume 4, EU Guidelines for Good Manufacturing Practice for Medicinal Products for Human and Veterinary Use, Annex 11: Computerised Systems.
- European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act).
- FDA. (2003). Guidance for Industry: Part 11, Electronic Records; Electronic Signatures, Scope and Application.
- FDA. (2004). Guidance for Industry: PAT, A Framework for Innovative Pharmaceutical Development, Manufacturing, and Quality Assurance.
- FDA. (2023). Using Artificial Intelligence and Machine Learning in the Development of Drug and Biological Products [Discussion paper]. U.S. Food and Drug Administration.
- FDA. (2011). Guidance for Industry: Process Validation, General Principles and Practices.
- FDA. (2018). Guidance for Industry: Data Integrity and Compliance With Drug CGMP, Questions and Answers.
- FDA (2021). Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan.
- MHRA. (2018). 'GxP' Data Integrity Guidance and Definitions (Revision 1).
- PIC/S. (2021). Good Practices for Data Management and Integrity in Regulated GMP/GDP Environments (PI 041-1).
- ISO/IEC (2023). ISO/IEC 23894: Artificial Intelligence — Risk Management.
- NIST (2022). Towards a Standard for Identifying and Managing Bias in Artificial Intelligence.

- **ICH and WHO lifecycle, tech transfer, and comparability (design to deployment)**

- ICH. (2009). ICH Q8(R2): Pharmaceutical Development.
- ICH. (2019). ICH Q12: Technical and Regulatory Considerations for Pharmaceutical Product Lifecycle Management.
- ICH. (2023). ICH Q9(R1): Quality Risk Management (Step 4).
- ICH. (2008). ICH Q10: Pharmaceutical Quality System.
- ICH. (2004). ICH Q5E: Comparability of Biotechnological/Biological Products Subject to Changes in Their Manufacturing Process.
- World Health Organization. (2011). WHO TRS 961, Annex 7: WHO Guidelines on Transfer of Technology in Pharmaceutical Manufacturing.
- World Health Organization. (2022). WHO TRS 1044, Annex 4: Technology Transfer in Pharmaceutical Manufacturing.
- Parenteral Drug Association. (2022). Technical Report No. 65 (Revised 2022): Technology Transfer. PDA.

- **Engineering, automation, and execution system standards (ISA-88/95 alignment, interoperability)**

- ASTM International. (2020). ASTM E2500-20: Standard Guide for Specification, Design, and Verification of Pharmaceutical and Biopharmaceutical Manufacturing Systems and Equipment.

- IEC. (2013 onwards, current edition varies by part). IEC 61512 (Batch Control), aligned to ISA-88.
- IEC. (2013 onwards, current edition varies by part). IEC 62264 (Enterprise-Control System Integration), aligned to ISA-95.
- OPC Foundation and IEC. (current consolidated). IEC 62541 (OPC UA).
- **Digital transformation and Pharma 4.0 framing**
- ISPE. (Pharma 4.0 initiative, current). Pharma 4.0 materials, operating model and holistic digital enablement resources.
- ISPE. (2022). GAMP 5, A Risk-Based Approach to Compliant GxP Computerized Systems (2nd Edition).
- ISPE (2023). Pharma 4.0™ Digital Maturity Model.
- ISPE (2019). Pharma 4.0™ Operating Model.
- World Economic Forum (2022). Global Lighthouse Network: Scaling Fourth Industrial Revolution Technologies in Manufacturing.
- Westerman, G., Bonnet, D., & McAfee, A. (2014). *Leading digital: Turning technology into business transformation*. Harvard Business Review Press.
- Malone, T. W. (2018). *Superminds: The surprising power of people and computers thinking together*. Little, Brown and Company.
- Vial, G. (2019). Understanding digital transformation: A review and a research agenda. *The Journal of Strategic Information Systems*, 28(2), 118-144. https://doi.org/10.1016/j.jsis.2019.01.003
- **AI governance, AI risk, and model lifecycle controls (relevant to Section 9 onward and governance Section 12)**
- ISO/IEC. (2023). ISO/IEC 42001: Artificial Intelligence Management System (AIMS), Requirements and Guidance.
- ISO/IEC. (2023). ISO/IEC 23894: Artificial Intelligence, Guidance on Risk Management.
- NIST. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1.
- ISPE. (2025). ISPE GAMP Guide: Artificial Intelligence.
- OECD (2021). OECD Framework for the Classification of AI Systems.
- OECD (2019). Recommendation of the Council on Artificial Intelligence.
- **AI enablement, data-centricity, and "AI-ready data" foundations (peer-reviewed anchors)**
- Jiménez-Luna, J., Grisoni, F. and Schneider, G. (2021). 'Artificial intelligence in drug discovery: recent advances and future perspectives'. RSC Chemical Biology.
- Sculley, D. et al. (2015). 'Hidden Technical Debt in Machine Learning Systems'. In: Advances in Neural Information Processing Systems (NeurIPS).
- Mitchell, M. et al. (2019). 'Model Cards for Model Reporting'. In: Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT).
- Gebru, T. et al. (2018). 'Datasheets for Datasets'. Communications of the ACM (preprint widely circulated in 2018; journal version published later).
- Wilkinson, M.D. et al. (2016). 'The FAIR Guiding Principles for scientific data management and stewardship'. Scientific Data.
- BioPhorum (2023). Data Enablement for AI in Biopharmaceutical Manufacturing.

- BioPhorum (2024). What Is AI-Ready Data? Concepts, Attributes, and Observability Considerations.
- World Economic Forum (2022). Blueprint for Scaling AI in Manufacturing.
- **Ontologies, semantics, and "industry ontology landscape" support**
- Pistoia Alliance (2023). CMC Process Ontology Project
- Pistoia Alliance (2024). CMC Process Ontology – Phase III Expansion
- NIIMBL. (Current). Big Data Program, Biopharmaceutical Manufacturing Ontology (project description and outputs).
- NIIMBL & Open Applications Group (OAGi) (2021). Biopharmaceutical Manufacturing Ontologies for Interoperability and Analytics. National Institute for Innovation in Manufacturing Biopharmaceuticals, Newark, DE.
- NIST. (2006). Ray, S.R. and Jones, A. 'Manufacturing Interoperability'. NIST publication.
- NIST. (2004). Ray, S.R. 'Tackling Semantic Interoperability at NIST'. NIST publication.
- NIST. (2004). Patil, L. et al. 'Ontology-Based Exchange of Product Data Semantics'. NIST publication.
- NIST. (2022). Kulvatunyou, B. et al. 'The Industrial Ontologies Foundry (IOF) Core Ontology'. NIST publication.
- CESMII. (Current). Smart Manufacturing Profiles (SM Profiles) and standards-based information models resources.
- Industrial Ontologies Foundry (IOF) (2022). Foundational Ontologies for Manufacturing and Industrial Systems.
- Arp, R., Smith, B., & Spear, A. D. (2015). Building Ontologies with Basic Formal Ontology. MIT Press.
- Conway, M. E. (1968). How do committees invent? Datamation, 14(4), 28–31.
- DAMA International. (2017). DAMA-DMBOK: Data Management Body of Knowledge (2nd ed.). Technics Publications.
- European Commission. (2022). EudraLex, Volume 4, EU Guidelines for Good Manufacturing Practice, Annex 16: Certification by a Qualified Person and Batch Release.
- Evans, E. (2003). Domain-Driven Design: Tackling Complexity in the Heart of Software. Addison-Wesley.
- Galbraith, J. R. (2014). Designing Organizations: Strategy, Structure, and Process at the Business Unit and Enterprise Levels (3rd ed.). Jossey-Bass.
- Kotter, J. P. (1996). Leading Change. Harvard Business School Press.
- Mintzberg, H. (1979). The Structuring of Organizations: A Synthesis of the Research. Prentice-Hall.
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature Machine Intelligence, 1(5), 206–215.
- Trist, E. L., & Bamforth, K. W. (1951). Some social and psychological consequences of the longwall method of coal-getting. Human Relations, 4(1), 3–38.
- Zachman, J. A. (1987). A framework for information systems architecture. *IBM Systems Journal*, 26(3), 276-292. https://doi.org/10.1147/sj.263.0276
- Ross, J. W., Weill, P., & Robertson, D. C. (2006). *Enterprise architecture as strategy: Creating a foundation for business execution*. Harvard Business School Press.
- Hohpe, G., & Woolf, B. (2003). *Enterprise integration patterns: Designing, building, and deploying messaging solutions*. Addison-Wesley.
- McKinsey & Company. (2016). *Closing the digital gap in pharma*. McKinsey & Company.

- Deloitte. (2021). *The biopharma digital transformation: Gain an edge with leapfrog digital innovation*. Deloitte Insights.
- BioPhorum. (2023). *Digital Plant Maturity Model (DPMM)*. BioPhorum Operations Group.
1. BioPhorum. (2026). *Managing Structured and Unstructured Data for AI in Pharma*. BioPhorum Operations Group.
- The Open Group. (2018). *The TOGAF® Standard, Version 9.2*. The Open Group.
- Yu, L. X. (2008). Pharmaceutical quality by design: Product and process development, understanding, and control. *Pharmaceutical Research*, 25(4), 781-791.
- Rathore, A. S., & Winkle, H. (2009). Quality by design for biopharmaceuticals. *Nature Biotechnology*, 27(1), 26-34.
- Garcia-Munoz, S., Dolph, S., & Ward, H. W. (2010). Handling uncertainty in the establishment of a design space for the manufacture of a pharmaceutical product. *Computers & Chemical Engineering*, 34(7), 1098-1107.
- Nasr, M. M., Krumme, M., Matsuda, Y., Trout, B. L., Badman, C., Mascia, S., ... & Lee, S. L. (2017). Regulatory perspectives on continuous pharmaceutical manufacturing: Moving from theory to practice. *Journal of Pharmaceutical Sciences*, 106(11), 3199-3206.
- Kourti, T. (2006). The Process Analytical Technology initiative and multivariate process analysis, monitoring and control. *Analytical and Bioanalytical Chemistry*, 384(5), 1043-1048.
- Peterson, J. J. (2008). A Bayesian approach to the ICH Q8 definition of design space. *Journal of Biopharmaceutical Statistics*, 18(5), 959-975.
- NIST. (2004). Ray, S. R. Interoperability standards in the semantic web. In *NIST Manufacturing Interoperability Program*.
- Hedberg, T. D., Lubell, J., Fischer, L., Maggiano, L., & Feeney, A. B. (2017). Testing the digital thread in support of model-based manufacturing and inspection. *Journal of Computing and Information Science in Engineering*, 16(2), 021001.
- Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In F.-J. Kahlen, S. Flumerfelt, & A. Alves (Eds.), *Transdisciplinary perspectives on complex systems* (pp. 85-113). Springer.
- IEC. (2010). IEC 61512-1: Batch control—Part 1: Models and terminology (3rd ed.). [ISA-88]
- IEC. (2013). IEC 62264-1: Enterprise-control system integration—Part 1: Models and terminology (2nd ed.). [ISA-95]
- Gruber, T. R. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2), 199-220.
- Fowler, M. (2017, February 7). What do you mean by "Event-Driven"? *martinfowler.com*. https://martinfowler.com/articles/201701-event-driven.html
- Hohpe, G. (2006, May). Programming without a call stack—Event-driven architectures. *ObjektSpektrum*.
- NIST. (2016). Measurement science roadmap for metal-based additive manufacturing. National Institute of Standards and Technology.
- Argote, L., & Ingram, P. (2000). Knowledge transfer: A basis for competitive advantage in firms. *Organizational Behavior and Human Decision Processes*, 82(1), 150-169.
- Kogut, B., & Zander, U. (1992). Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science*, 3(3), 383-397.
- ISPE. (2018). ISPE Good Practice Guide: Technology transfer (3rd ed.). International Society for Pharmaceutical Engineering. 8. BioPhorum. (2020). Technology transfer in biopharmaceutical manufacturing: Best practices and recommendations.
- Farid, S. S. (2007). Process economics of industrial monoclonal antibody manufacture. *Journal of Chromatography B*, 848(1), 8-18.

- Garcia-Aponte, O. F., Herwig, C., & Gortez, B. (2017). Towards integrated online monitoring of cell culture processes. *Biotechnology Progress*, 33(6), 1433-1445.

-

-

- **Source links (for traceability)**

- EU AI Act (Regulation (EU) 2024/1689): https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng
FDA Process Validation (2011 PDF): https://www.fda.gov/files/drugs/published/Process-Validation--General-Principles-and-Practices.pdf
FDA Data Integrity Q&A (2018 PDF): https://www.fda.gov/media/119267/download
21 CFR Part 11 (eCFR): https://www.ecfr.gov/current/title-21/chapter-I/subchapter-A/part-11
EU GMP Annex 11 (2011 PDF): https://health.ec.europa.eu/system/files/2016-11/annex11_01-2011_en_0.pdf
MHRA GxP Data Integrity (2018 PDF): https://assets.publishing.service.gov.uk/media/5aa2b9ede5274a3e391e37f3/MHRA_GxP_data_integrity_guide_March_edited_Final.pdf
PIC/S PI 041-1 (landing, PDF accessible): https://picscheme.org/docview/4234
ICH Q8(R2) PDF: https://database.ich.org/sites/default/files/Q8_R2_Guideline.pdf
ICH Q10 PDF: https://database.ich.org/sites/default/files/Q10%20Guideline.pdf
ICH Q12 PDF: https://database.ich.org/sites/default/files/Q12_Guideline_Step4_2019_1119.pdf
ICH Q9(R1) (Step 4 PDF): https://database.ich.org/sites/default/files/ICH_Q9%28R1%29_Guideline_Step4_2025_0115.pdf
ICH Q5E PDF: https://database.ich.org/sites/default/files/Q5E%20Guideline.pdf
WHO TRS 961 Annex 7 (Tech transfer, 2011 PDF): https://www.who.int/docs/default-source/medicines/norms-and-standards/guidelines/production/trs961-annex7-transfer-technology-pharmaceutical-manufacturing.pdf
WHO TRS 1044 Annex 4 (Tech transfer, 2022 PDF): https://cdn.who.int/media/docs/default-source/medicines/norms-and-standards/guidelines/production/trs1044-annex4-technology-transfer-in-pharmaceutical-manufacturing.pdf
FDA PAT Guidance (2004 PDF): https://www.fda.gov/media/71012/download
NIST AI RMF 1.0 PDF: https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf
ISO/IEC 42001 summary page: https://www.iso.org/standard/42001
ISO/IEC 23894 summary page: https://www.iso.org/standard/77304.html
Pistoia Alliance CMC Ontology project: https://pistoiaalliance.org/project/cmc-ontology/
NIIMBL Biopharmaceutical Manufacturing Ontology: https://www.niimbl.org/projects/niimbl-big-data-program-biopharmaceutical-manufacturing-ontology/
CESMII Smart Manufacturing Profiles: https://www.cesmii.org/technology/sm-profiles/
ASTM E2500-20 landing: https://www.astm.org/e2500-20.html
ISPE Pharma 4.0 initiative: https://ispe.org/initiatives/pharma-4.0
ISPE GAMP 5 2nd Edition landing: https://ispe.org/publications/guidance-documents/gamp-5-guide-2nd-edition
PubMed entry (AI in drug discovery review): https://pubmed.ncbi.nlm.nih.gov/33779453/

-
-

• APPENDIX A: ALARM MANAGEMENT PLATFORM CASE STUDY

1. EVIDENCE GENERATION THROUGH ARCHITECTURAL DESIGN: A PRACTITIONER IMPLEMENTATION

- **Context:** Cell therapy manufacturing organization, clinical-stage biopharmaceutical production

1.1. PROBLEM STATE: THE CHALLENGE OF DISTRIBUTED ALARM MANAGEMENT

- **The Scale of Distributed Alarm Data.** Cell therapy manufacturing environments generate alarms across diverse equipment types—bioreactors, incubators, refrigerators, freezers, centrifuges, and analytical instruments. In the absence of a centralized platform, alarm data accumulates across individual systems, each with its own interface, storage format, and retention approach. This distribution is not a failure of design but a natural consequence of equipment being procured and deployed to meet evolving operational needs. The challenge becomes apparent when organizations seek to review alarm history holistically: recovery requires accessing each piece of equipment individually, and in many cases historical data has been archived to file shares where retrieval is neither straightforward nor timely.
- **The Temporal Gap in Alarm Review.** When alarm acknowledgment and quality review occur separately from alarm generation—sometimes by days or weeks—organizations face an inherent challenge in reconstructing context. Alarms logged in batch records by operators were typically reviewed only when a deviation triggered an investigation. This approach is rational given resource constraints: not every alarm warrants detailed review, and quality organizations must prioritize their attention. However, it creates temporal disconnect between when an alarm occurred, when it was documented, and when its significance was assessed. This pattern reflects a broader industry challenge rather than an organizational shortcoming. The episodic, post-hoc evidence assembly described in Section 6.3 is common across pharmaceutical manufacturing precisely because it evolved within regulatory frameworks that emphasized final product testing over continuous process verification.
- **The Stakes in Cell Therapy Manufacturing.** The consequences of delayed alarm visibility are amplified in cell therapy contexts. In autologous therapies, where one batch represents one patient, an alarm that goes unnoticed during a critical process window may contribute to batch loss. Such losses extend beyond operational metrics—they represent patients who may need to repeat the substantial physical and emotional burden of providing starting material for their own therapy. This patient-centric framing underscores why architectural approaches to alarm management warrant investment, even when existing processes are functioning adequately by traditional measures.
- **Opportunity for Architectural Improvement.** Quality Operations recognized an opportunity to move from reactive alarm review to proactive disposition. The personnel who would benefit most from a centralized platform had been seeking solutions but were constrained by technical complexity and competing priorities. The implementation described here emerged from this recognition—not as a correction of failure, but as an architectural evolution enabling capabilities that distributed systems could not practically provide.

1.2. Solution Architecture

- The Alarm Management Platform (AMP) was built on commercial SCADA infrastructure, utilizing edge nodes to capture and communicate data with equipment, processed by backend gateways, and presented via web interface. This was a purpose-built solution—not an off-the-shelf alarm management product—designed around three core capabilities: (1) Unified Presentation: present alarms from heterogeneous equipment sources in a single interface, eliminating the need to access each system individually; (2) Operator Interaction: enable operators to review, comment on, and acknowledge alarms with contextual annotation at or near the time of occurrence; and (3) Quality Review: allow quality reviewers to view and disposition alarms in real time alongside operations, shifting from retrospective reconstruction to contemporaneous assessment.

1.3. IDDC Architectural Mapping

- The AMP implementation provides a concrete instantiation of three IDDC layers, demonstrating how the architectural framework translates into operational capability.
- **Event Layer.** The Event Layer captures alarms and their relationships to equipment, process, and batch. Each alarm is stored as a discrete, time-stamped event with equipment identifier and role, process context (batch, phase, step), alarm type and severity classification, timestamp with sub-second precision, and initial state and triggering conditions. This event-centric representation enables accurate reconstruction of what occurred without retrospective inference—addressing the temporal coherence requirements described in Section 6.4.
- **Context Layer.** The Context Layer binds human interpretation from operators and quality reviewers to each alarm event. This contextualization occurs contemporaneously with execution, documenting causation (why the alarm occurred), treatment (how it was addressed), resolution (whether and how it was resolved), and impact assessment (effect on batch quality and regulatory status). This binding of human judgment to execution events exemplifies what Malone (2018) characterizes as cyber-human collaboration essential to collective intelligence in socio-technical systems. The operator provides domain expertise and situational awareness; the system provides comprehensive visibility and structured capture. Together, they produce contextualized evidence that neither could generate alone.
- **Evidence Layer.** The Evidence Layer produces structured output suitable for quality review and regulatory inspection: alarm lineage (complete history from generation through disposition), audit trail (personnel involvement with timestamps and actions), decision framework (rationale supporting batch status determination), and regulatory alignment (evidence structured to satisfy ALCOA+ principles). This evidence builds up during batch execution rather than being assembled after the fact—enabling the continuous disposition model described in Section 8.
- **Semantic Layer Gap.** The current implementation does not fully realize the Semantic Layer. The intent of each alarm—why it exists, what CPP or CQA it protects—remains tacit knowledge rather than formally encoded in an ontological structure. A full IDDC implementation would formalize this alarm intent, enabling automated reasoning about alarm significance and supporting more sophisticated AI-assisted disposition (see Section 10).

1.4. Change Management Approach

- **Strategic Influencer Model.** Rather than approaching senior decision-makers directly with a technology proposal, the implementation team engaged a dedicated change manager who identified the right influencer within each organization—the person who understood operational realities and could frame the approach constructively for their peers and leadership. This reflects the principle that collective intelligence depends not only on individual capabilities but on the communication structures that connect them.
- **Co-Development Partnership.** Quality Operations provided strong support once the platform's potential became clear. Rather than receiving a finished tool, they became partners in development—contributing requirements, validating design decisions, and shaping the solution they would ultimately use. Specific contributions included definition of alarm disposition states, specification of critical information requirements, workflow design for operator-quality reviewer interaction, and acceptance criteria for evidence completeness. This partnership model improved solution fitness and created ownership that accelerated adoption.

1.5. Outcomes and Observed Impact

- Before implementation, preparing batch-related alarm data for disposition was a retrospective, multi-step reconstruction: identifying the equipment involved in a batch, establishing the time windows during which that equipment was in use, retrieving the alarms and events generated on that equipment during those windows, and delivering the assembled record to manufacturing and then quality reviewers only after manufacturing disposition had completed. This reconstruction typically took five to ten days, and extended further whenever an alarm required investigation or opened a deviation.
- With the AMP, alarms are bound not merely to a batch but to the specific batch phase on the specific equipment, allowing alarms outside relevant operational phases to be excluded and critical-window alarms to be elevated concisely. Meaning is assigned to an alarm while it is still active and being triaged—or shortly after—by the operator or technician handling it, producing contemporaneous, structured input whose quality is higher because causation, response, resolution, and impact are recorded when they are clearest rather than reconstructed weeks later. Operators return to other work while disposition evolves, while quality representatives review the initial acknowledgement and actions in parallel, assess impact in real time, and draw on previously dispositioned similar or identical alarm situations to disposition alarms ahead of, or shortly after, batch completion.
- Having all batch-related alarms dispositioned and review-ready reduced the disposition step by an average of two to three days. Measured against time-to-close-and-release, the overall saving was approximately three to five days per batch. Where an alarm still required a deviation, the investigation itself was shortened by one to two days, because operators no longer had to be re-engaged to reconstruct actions and impact for an alarm that had occurred days or weeks earlier. The platform did not remove the need to triage an active alarm and therefore did not materially change deviation detection; the savings were concentrated in post-batch manufacturing and quality review and in investigation. Incorporating this dispositioned alarm data and its disposition state into the electronic batch record for review by exception is expected to compress disposition time further.
- Other effects were observed but not formally quantified. Improved alarm rationalization and classification, together with stronger notification and escalation, focused response and surfaced nuisance alarms for elimination, reducing the overall alarm burden; per-alarm

review effort itself changed little, the benefit being the ability to begin review before batch completion. Reviewer hours and alarm-volume reductions were not measured, and these results reflect a single-site implementation rather than a controlled study—consistent with the evaluation scope described in the paper's Limitations.

1.6. AI Enablement Pathway

- The platform's digital, extensible architecture enables future use of AI models to pre-disposition certain alarms within defined boundaries, further reducing burden on quality personnel. This represents an entry point to the agentic AI capabilities discussed in Section 10—demonstrating how IDDC infrastructure compounds in value over time as more sophisticated capabilities are layered upon the architectural foundation. Potential AI-assisted capabilities include pattern recognition (identifying alarm patterns associated with known root causes), risk stratification (prioritizing alarms requiring human attention based on historical impact), pre-disposition (automated disposition of routine, well-characterized alarm types within validated boundaries), and anomaly detection (flagging unusual alarm patterns that deviate from learned baselines). Each capability requires the semantic, event, context, and evidence structures that the AMP provides—validating the IDDC thesis that architectural readiness precedes AI enablement.

1.7. Lessons for IDDC Implementation

- The AMP case study offers several lessons applicable to IDDC implementations more broadly. First, narrow architectural focus enables rapid delivery of demonstrable value while establishing patterns for future extension. Second, cyber-human design—explicitly positioning humans and machines in complementary roles—characterizes effective implementations. Third, treating evidence generation as a first-class architectural requirement distinguishes implementations that enable continuous disposition from those that merely digitize retrospective assembly. Finally, technical architecture alone is insufficient; change management, strategic influencer engagement, and co-development partnerships are essential to translating architectural capability into operational value.
- *Note: Organization and platform vendor names have been anonymized. The author served as solution architect, sponsor, and implementer for this platform.*

• APPENDIX B: ONTOLOGY AS ARCHITECTURAL INFRASTRUCTURE IN MANUFACTURING

1. FROM SEMANTIC THEORY TO DEPLOYMENT REALITY: A PRACTITIONER ARCHITECTURE

- **Context:** Cell and gene therapy manufacturing as bounded design scope; biopharmaceutical ontology ecosystem as architectural foundation
- **Framing:** This appendix presents a design science artifact—a proposed ontology deployment architecture validated by alignment with international standards, consortium consensus (Pistoia Alliance, IOF, NIIMBL), and ISA-95/ISA-88 structural models. It addresses a question that existing ontology literature largely leaves unanswered: what does deploying an ontology in a manufacturing facility actually require of engineers, architects, and the systems they govern?

1.1. The Deployment Question

1.1.1. WHAT ONTOLOGY LITERATURE DOES NOT ANSWER

- The ontology literature for life sciences manufacturing has matured substantially. Foundational work on Basic Formal Ontology (Arp, Smith, & Spear, 2015), the Industrial Ontologies Foundry (IOF) Core specification, and domain-specific efforts such as the Pistoia Alliance CMC Process Ontology (CMC-O) and NIIMBL Biopharmaceutical Manufacturing Ontology (BPMO) have established a layered semantic framework spanning top-level philosophical distinctions through domain-specific manufacturing constructs.
- What this literature does not systematically address is the architectural deployment question: where does the ontology reside in a manufacturing facility? How do operational systems—MES, SCADA, LIMS, historians—connect to and consume ontological structures? Who governs semantic definitions, and through what mechanisms do changes propagate across consuming systems? These are not theoretical concerns. They determine whether an ontology functions as an architectural enabler or remains an academic artifact disconnected from operational reality.

1.1.2. THE DISTINCTION BETWEEN DEFINING MEANING AND DEPLOYING IT

- An ontology defines entities, roles, relationships, and constraints that give data its meaning within a domain. Deploying that ontology in manufacturing requires translating these definitions into structures that operational systems can instantiate, integration layers can transport, data platforms can persist, and quality processes can consume. This translation is neither automatic nor trivial—it demands deliberate architectural design across multiple system boundaries.
- The architectural argument of this appendix is therefore: **an ontology functions as a governed semantic backbone defining meaning, roles, and relationships—not as an operational system. Systems map to it rather than execute it.** Relationships are defined centrally, instantiated in MES, mapped at integration boundaries, materialized in data platforms, and consumed by analytics and quality processes.

1.2. The Biopharmaceutical Manufacturing Ontology Stack

1.2.1. LAYERED ARCHITECTURE

- The biopharmaceutical manufacturing ontology ecosystem is converging toward a layered, interoperable architecture (see Figure 4 in the main text). This layering reflects a deliberate design principle: domain-agnostic foundations enable cross-industry reasoning, while progressively specialized layers address pharmaceutical, manufacturing, and modality-specific requirements.
- **Foundational Layer (Domain-Agnostic):** Basic Formal Ontology (BFO) serves as the ISO-recognized top-level ontology (ISO/IEC 21838-2:2020), providing distinctions between continuants and occurrents, roles, qualities, and processes that are invariant across domains. BFO's formal rigor ensures that downstream ontologies inherit consistent logical structure, enabling reasoning that transcends any single manufacturing context.
- **Mid-Level Layer (Cross-Industry Manufacturing):** The Industrial Ontologies Foundry (IOF) Core extends BFO into manufacturing, defining reusable patterns for planned processes, material roles, design specifications, events, and assets. IOF Core provides the semantic bridge between domain-agnostic philosophy and domain-specific manufacturing reality—establishing, for example, how a "planned process" relates to its realization, how materials participate in processes through defined roles, and how information content entities relate to the physical entities they describe.
- **Domain Layer (Pharmaceutical and Regulatory):** Multiple concurrent efforts address pharmaceutical-specific semantics:
- **CMC-O:** Pistoia Alliance CMC Ontology (CMC-O): Focused on ISA-88 and ISA-95 aligned process, recipe, and parameter semantics. CMC-O defines the Process Stage → Operation → Action hierarchy that structures how manufacturing processes are described, compared, and transferred.
- **IDMP-O:** Operationalizes ISO IDMP standards into a computable ontology for medicinal product identity, composition, packaging, and regulatory context—providing the product-level semantics that link manufacturing execution to regulatory filings.
- **NIIMBL BPMO:** Adds biologics-specific unit operation semantics at a finer-grained operational level than CMC-O, addressing the particular requirements of upstream and downstream bioprocessing.
- **Allotrope Foundation Ontologies (AFO):** Standardize laboratory instrument, method, and result semantics, bridging manufacturing execution data with analytical and QC data.
- **QUDT:** Provides canonical semantics for quantities, units, and dimensions—ensuring measurement consistency across laboratory, manufacturing, and data platform boundaries.
- **Provenance Layer:** The W3C Provenance Ontology (PROV-O) offers a general model for representing data lineage, activities, agents, and derivations. When profiled for GMP use, PROV-O supports the audit trail, electronic signature, and traceability requirements that are foundational to batch disposition and regulatory inspection.

1.2.2. MATURITY ASSESSMENT

- The ontology ecosystem spans a range of maturity states, which has direct implications for deployment architecture:

- **Established and production-ready:** BFO, OBO Foundry ontologies, IDMP-O (v1+), Allotrope Foundation Ontologies, QUDT.
- **Released but maturing:** IOF Core (v1 beta, expanding adoption), Pistoia Pharma General Ontology (Phase 1 published).
- **Actively in progress:** Pistoia CMC-O, NIIMBL BPMO.
- This heterogeneous maturity means that a deployment architecture cannot assume all semantic layers are stable. The architecture must accommodate ontologies at different maturity stages—some as authoritative references, others as evolving consensus artifacts subject to versioned change management.

1.2.3. IDENTIFIED GAPS

- Four gaps are architecturally significant for manufacturing deployment:

a) **Digital artifact semantics.** No widely adopted ontology module represents digital artifacts such as control logic, models, algorithms, and their lifecycle in GMP manufacturing. As AI and advanced analytics become execution dependencies, this gap becomes increasingly consequential.
b) **Measurement modeling consistency.** Despite QUDT's canonical representation of units, consistent patterns for representing measured qualities, specification limits, and measurement uncertainty across laboratory and plant data remain unevenly implemented.
c) **Cell and gene therapy semantics.** Chain-of-identity, chain-of-custody, and patient linkage concepts are only partially addressed and lack a unified reference profile—a gap that becomes critical when ontology is deployed in CGT manufacturing contexts (see the CGT Manufacturing stress-test discussion later in this appendix).
d) **Batch disposition domain model.** No open reference ontology module explicitly models batch disposition decisions, release-by-exception logic, and QA review artifacts in alignment with manufacturing, quality, and regulatory ontologies. Batch disposition remains an emergent capability that must be composed from multiple semantic layers rather than instantiated from a coherent disposition-specific model.

1.3. Deployment Architecture: Six Layers of Ontological Mapping

- The central architectural contribution of this appendix is a six-layer model describing where and how ontological relationships are realized in a manufacturing facility. Each layer serves a distinct function; together, they translate governed semantic definitions into operational capability.

1.3.1. GOVERNED ONTOLOGY LAYER

- The governed ontology is not an operational system. It is the authoritative semantic reference to which all downstream systems map. In IOF Core / BFO terms, this layer defines role patterns, participation patterns, aboutness relationships, and plan and protocol specification semantics. Governance of this layer requires versioning, change control, and clear ownership—the same disciplines applied to other lifecycle-critical assets such as master batch records and regulatory filings.

1.3.2. ONTOLOGY ACCESS AND DISTRIBUTION: HOW SYSTEMS STAY CURRENT

- The deployment question that most directly affects practicing engineers is deceptively simple: where does the ontology actually live, and how do the systems that depend on semantic consistency—MES, SCADA, PLM, historians, LIMS, integration layers, data

platforms, digital tech transfer platforms, and quality systems—access it? If this question is not answered architecturally, the ontology remains an intellectual asset rather than an operational one.

- **Persistence and Hosting.** The governed ontology is maintained as a formal artifact—typically serialized in OWL (Web Ontology Language) or a functionally equivalent format—within a dedicated ontology management platform or repository. This is not a static file on a shared drive. It is a versioned, access-controlled asset served through an application layer that supports querying, validation, and change management. Candidate hosting approaches include dedicated ontology servers (e.g., platforms supporting SPARQL endpoints), semantic middleware embedded within integration infrastructure, or purpose-built ontology registries that expose semantic definitions through APIs. The specific technology choice is less architecturally significant than the requirement that the ontology be accessible programmatically, version-controlled, and governed with the same rigor as other lifecycle-critical reference data.

- **Programmatic Access as the Architectural Principle.** The defining characteristic of an externalized ontology—what distinguishes it from embedded semantics—is that consuming systems access semantic definitions programmatically rather than through manual transcription. If engineers must read the ontology, interpret it, and manually configure each system to reflect its content, the architecture has not solved the harmonization problem; it has merely centralized the specification while leaving the same manual, error-prone translation in place. The architectural target is systems that consume the ontology directly through APIs, validate their configurations against it automatically, and synchronize when it changes—with human oversight governing the change process, not performing the transcription.

- **Access Model.** The ontology repository exposes semantic definitions through queryable interfaces (SPARQL endpoints, REST APIs, or equivalent) that any consuming system can call. This access model applies uniformly across the architecture. Every system that creates, transforms, transports, or evaluates manufacturing data is a potential consumer, and all of them consume the same authoritative source through the same programmatic mechanisms:

- **Configuration validation:** When PLM specifications are structured, MES master data is configured, SCADA alarm definitions are created, historian tag structures are defined, or UNS topic hierarchies are established, the system queries the ontology repository to validate that local definitions align with the current governed version. Misalignments are flagged automatically rather than discovered during downstream reconciliation or audit. This is particularly consequential at the PLM-to-MES boundary, where process specifications—limits, setpoints, material definitions, quality attributes—are currently transcribed manually from documents into execution system configurations. When both PLM and MES consume the same ontology programmatically, the specification and its executable instantiation reference the same governed semantic definitions, eliminating the transcription step and the validation burden it creates.

- **Synchronized updates:** When the ontology evolves—whether through new process understanding, regulatory feedback, or consortium updates—the repository notifies consuming systems of the change. Each system assesses impact against its current configuration, identifies definitions that require update, and presents the required changes for validated deployment. The change propagation is automated; the approval and deployment follow established change control procedures.

- **Runtime resolution:** During execution, systems resolve ontology references through API calls or cached local replicas synchronized with the authoritative source. An alarm generated by SCADA can be contextualized against the current ontology to determine which CPP or CQA it protects. An MES batch record can reference the ontology to verify that material roles and process relationships are semantically current. A historian query can be enriched with ontological context to ensure that data retrieved from different time periods or different equipment is interpreted consistently. A digital tech transfer platform can compare process definitions across sites by resolving both against the same ontology, identifying semantic equivalences and gaps without manual cross-referencing. A PLM system can verify that a specification change is reflected consistently across every downstream system that consumes it. Runtime resolution extends to any system that benefits from semantic context, which in a manufacturing facility is effectively all of them.
- **Audit and compliance queries:** Quality, regulatory, and governance functions query the ontology repository to verify that operational systems are aligned with the current semantic version, to trace how semantic definitions have evolved over time, and to assess the impact of proposed changes before they propagate. Because every consuming system declares which ontology version it implements, compliance verification becomes a queryable property of the architecture rather than a manual audit exercise.
- **Version Synchronization.** Because multiple systems consume the same ontology, version management is architecturally critical. The ontology repository maintains explicit version history with change documentation. Consuming systems declare which semantic version they implement, creating a traceable map of semantic alignment across the facility. When a new version is published, the propagation follows a managed process: impact assessment is automated, configuration updates are generated programmatically, validation occurs against the new version, and deployment follows established change control.
- **The Alternative: Embedded Semantics.** Without this access architecture, semantic definitions are embedded directly into each consuming system—authored in PLM documents, manually transcribed into MES master data, programmed into SCADA alarm descriptions, replicated across historian tag structures, duplicated in integration mappings, and reconstructed independently in tech transfer documentation for receiving sites. This approach works, and it is the current industry norm. Its cost is measured in harmonization effort and transcription risk: when a specification changes in PLM, every downstream system that consumed that specification must be identified, assessed, and updated through manual processes that are themselves sources of error. When a process transfers to a new site or CDMO, the semantic definitions embedded in the sending site's systems must be extracted, interpreted, and re-embedded into the receiving site's systems—precisely the lifecycle discontinuity described in Section 4. The ontology access architecture described here replaces that manual harmonization with programmatic synchronization—systems that consume a shared, versioned semantic reference and stay aligned through automated validation rather than human transcription.

1.3.3. MES / EBR INSTANTIATION LAYER (ISA-95 LEVEL 3)

- At batch creation, the Manufacturing Execution System (MES) instantiates a network of ontological relationships, drawing on the governed ontology accessed through the mechanisms described in the Ontology Access and Distribution subsection of this appendix. ISA-95 structural objects provide the execution model; the ontology supplies meaning, roles, and traceability. Product Definitions map to Plan Specifications, Production Orders to Planned Processes, Material Lots to Material Entities with roles, Equipment to Material

Artifacts, Personnel to Agents, and Production Schedules to Temporal Regions. The MES does not "run" the ontology—it instantiates relationships that the ontology defines, creating traceable links between manufacturing intent and execution reality.

1.3.4. INTEGRATION / UNIFIED NAMESPACE LAYER (ISA-95 LEVEL 3–2 BOUNDARY)

- At the integration boundary, ontological mappings contextualize execution data as it flows between systems: measurements are linked to the material lots they describe, material lots are linked to their derivation history, events are recognized within the context of specific processes and batches, and equipment and personnel participation is recorded with temporal precision. The integration layer—whether implemented as a Unified Namespace (UNS), event bus, or traditional middleware—must preserve these semantic relationships during transport.

1.3.5. DATA PLATFORM / KNOWLEDGE GRAPH LAYER (ISA-95 LEVEL 4)

- The data platform materializes ontological relationships as persistent, queryable structures. This layer enables capabilities that operational systems cannot practically provide: cross-batch comparison, cross-site equivalence assessment, trend analysis grounded in semantic consistency, and the knowledge graph structures that support advanced analytics and AI. Ontology identifiers serve as lineage anchors, cross-site equivalence mappings support technology transfer, and temporal versioning of semantic definitions preserves the ability to evaluate historical data against the intent in effect at the time of manufacture.

1.3.6. CONSUMPTION LAYER: DISPOSITION, QMS, SUPPLY CHAIN, AND AI

- The consumption layer is where the architectural investment in ontology produces its most visible returns. Batch disposition is not a standalone function applied to data after the fact. It is an emergent capability that depends on the consistent accumulation of semantically grounded evidence across all upstream layers. When the ontology architecture is intact, disposition becomes a role assertion on a material entity supported by traceable relationships spanning process execution, data generation, deviations, and specification conformance—the continuous disposition model described in Section 8.

1.4. CGT Manufacturing as Architectural Stress Test

- Cell and gene therapy (CGT) manufacturing was selected as the bounded scope for this architectural analysis because it intensifies every dimension of the ontology deployment problem. Where small molecule manufacturing permits statistical abstraction, pooled lots, and aggregated quality assessment, CGT manufacturing demands: mandatory lot genealogy where no statistical sampling or lot aggregation is permitted; patient linkage where material identity is inseparable from the patient it serves; non-collapsible lineage where every derivation is individually traceable; and patient-specific disposition semantics where batch release applies to a therapy intended for a specific human being.

- Each layer of the deployment architecture is intensified in CGT contexts. The governed ontology must include patient-linking role patterns and chain-of-identity semantics. MES production orders become patient-specific. Integration layer lot genealogy is non-collapsible. Contextualization services must maintain bidirectional genealogy across processing steps spanning days to months. The data platform must persist patient entity relationships across the full lifecycle including CDMO partnerships. And disposition applies to a specific patient's therapy rather than a statistical population.

- CGT manufacturing thus serves as a stress test that reveals architectural requirements which remain latent in less demanding modalities. The ontology deployment architecture that satisfies CGT requirements inherently satisfies small molecule requirements—but the reverse is not true.

1.5. ISA-95 as Structure, Ontology as Meaning

- A persistent concern in manufacturing digitalization is whether ontology-based architectures conflict with established automation standards. The architecture presented here is explicitly designed to demonstrate that ISA-95 and ontology serve complementary, non-competing functions. ISA-95 provides structural execution models: the object types, hierarchies, and information flows that automation systems implement. Ontology provides semantic meaning: the roles, relationships, and traceability that give ISA-95 objects their interpretive context. A Material Lot in ISA-95 terms is a structural container; in ontological terms, it is a Material Entity bearing specific roles with defined participation in processes and traceable derivation history.
- This complementarity means ontology deployment does not require replacing ISA-95 infrastructure. It requires extending existing structural objects with semantic annotations that enable cross-system interpretation, lifecycle traceability, and knowledge preservation. Manufacturing facilities operate within validated ISA-95 infrastructure that regulators understand and inspect against. An ontology deployment architecture that respects ISA-95 structure while adding semantic capabilities can be presented to regulators as an enhancement to existing validated systems, not a replacement of them.

1.6. Current State and Honest Assessment

1.6.1. INDUSTRY MATURITY

- The biopharmaceutical industry's relationship with manufacturing ontology is characterized by a widening gap between semantic infrastructure readiness and operational deployment. On the infrastructure side, substantial progress is evident: BFO is ISO-standardized, IOF Core is expanding adoption across manufacturing sectors, the Pistoia Alliance CMC-O and NIIMBL BPMO are actively developing domain-specific constructs, and digital infrastructure prerequisites—MQTT-capable edge architectures, modern SCADA platforms, cloud-connected MES systems—are increasingly common in new and renovated facilities.
- On the deployment side, however, no published example of an end-to-end ontology deployment architecture in biopharmaceutical manufacturing has been identified by the author. Research and laboratory ontology implementations exist at several organizations, and individual ontology layers (particularly IDMP-O for regulatory submissions and AFO for analytical data) are in production use. But the full architectural stack described in this appendix—governed ontology through MES instantiation through integration-layer contextualization through knowledge graph materialization through quality consumption—remains unrealized at industry scale.
- This gap is not primarily technical. The ontology components exist or are maturing; the standards alignment is achievable; the architectural patterns are sound. What remains is the organizational commitment, cross-functional coordination, and sustained investment required to compose these elements into a functioning deployment. The design science contribution of this appendix is the architecture itself—providing a validated blueprint that reduces the engineering uncertainty organizations face when making that investment.

1.6.2. CROSS-INDUSTRY PERSPECTIVE

- Other industries offer partial precedent. ISO 15926 in process industries (oil and gas) provides a long-running example of ontology-based information integration, though its scope and adoption challenges offer both encouragement and cautionary lessons. The Digital Twin Definition Language (DTDL) in discrete manufacturing demonstrates that semantic metadata can be bound to operational assets at scale. These cross-industry examples validate the general architectural pattern while highlighting that biopharmaceutical manufacturing's unique regulatory, quality, and patient-safety requirements demand domain-specific instantiation.

1.7. IMPLICATIONS FOR THE IDDC FRAMEWORK

- The ontology deployment architecture described here directly supports the IDDC's Semantic Layer—the foundational layer that strengthens the consistency and scalability of Event, Context, and Evidence capabilities. It is important to be precise about what ontology adds and what already functions without it. Manufacturing organizations routinely program semantic context into their systems—alarm descriptions reference process parameters, MES recipes encode relationships between materials and equipment, and integration layers map data between systems. These implementations work. The architectural concern is not that context is absent without ontology, but that it is implemented redundantly, maintained independently, and harmonized manually across systems and lifecycle phases:

- **Events:** Events are captured with context programmed into each generating system—but that context is defined locally, creating risk of inconsistency when the same concept (e.g., a critical process parameter, an equipment role) is represented differently across SCADA, MES, and historian systems. Reconciling these representations requires manual intervention at each integration boundary.

- **Context:** Context is added by humans and by system configuration—but without a shared semantic reference, maintaining consistency across systems demands ongoing harmonization effort that scales poorly as facilities grow in complexity. Two systems may both contextualize an alarm correctly in isolation while producing irreconcilable interpretations when their data converges downstream.

- **Evidence:** Evidence is assembled through established quality processes—but the manual effort required to verify that evidence from multiple sources is semantically consistent represents a significant and recurring cost, particularly during batch disposition, technology transfer, and regulatory inspection.

- Ontology does not create these capabilities where none exist. It externalizes and governs the semantic definitions that are currently embedded, duplicated, and independently maintained across systems—reducing harmonization burden, improving consistency at scale, and enabling the progressive evidence accumulation that supports continuous disposition (Section 8).

- The AMP case study (Appendix A) illustrates this distinction concretely: the platform successfully realizes Event, Context, and Evidence layers through programmed context and human interaction. The Semantic Layer—formal encoding of alarm intent in ontological structures—remains unrealized, meaning the platform's semantic definitions are maintained within the application rather than governed externally. The ontology deployment architecture described here would externalize those definitions, enabling

consistent interpretation across AMP and all other systems that consume the same manufacturing concepts.

- *Note: The author is a contributing member of the Pistoia Alliance CMC Process Ontology project. The architectural analysis presented here draws on publicly available consortium outputs, international standards, and the author's practitioner experience in cell therapy manufacturing digitalization. No proprietary organizational data is disclosed.*

• APPENDIX C: RESEARCH METHODOLOGY

2. METHODOLOGICAL FRAMEWORK

- This research adopts a design science research methodology (Hevner et al., 2004; Peffers et al., 2007) to develop and evaluate an architectural framework for lifecycle continuity in pharmaceutical manufacturing. Design science is particularly appropriate for this work because it addresses a class of problems—digital transformation failures in regulated manufacturing—through the creation and evaluation of innovative artifacts (March & Smith, 1995). The primary artifact produced by this research is the Intent-Driven Digital Core: a reference architecture framework that prescribes how manufacturing intent, execution context, and quality evidence should be preserved across the product lifecycle.
- Unlike purely empirical research that seeks to understand phenomena through observation, or purely theoretical research that develops conceptual models without implementation consideration, design science research bridges theory and practice by creating utility-oriented solutions grounded in rigorous analysis (van Aken, 2004; Gregor & Hevner, 2013). This methodological positioning acknowledges that the value of the proposed framework lies not only in its conceptual coherence but in its practical applicability to real-world manufacturing challenges.

3. RESEARCH ACTIVITIES

- The research followed the design science methodology proposed by Peffers et al. (2007), which structures design science research into six interconnected activities: problem identification and motivation, objectives of solution, design and development, demonstration, evaluation, and communication.

3.1. PROBLEM IDENTIFICATION AND MOTIVATION

- The research problem emerged from extensive engagement with pharmaceutical manufacturing organizations, grounded in over a decade of professional experience in pharmaceutical and biotechnology manufacturing systems and process automation. This practical foundation was complemented by active industry consortium participation, particularly the Pistoia Alliance CMC Process Ontology project (2023-2025) and review of BioPhorum Data Enablement Program materials. Further academic grounding is provided through the author's pursuit of the MIT Sloan Credential in Digital Business, including coursework in AI strategy, creative transformation, and algorithmic business thinking—perspectives that inform the architectural approach developed in this work. Analysis of digital transformation initiatives across the industry revealed a consistent pattern: substantial investment in digital technologies accompanied by limited realization of sustained, scalable value. Initial investigation focused on identifying whether these failures

stemmed from technological inadequacy, organizational resistance, or more fundamental structural issues.

- Through systematic analysis of technology transfer documentation, batch disposition processes, and digital maturity assessments, architectural discontinuity emerged as a primary causal factor. Manufacturing intent, execution context, and quality evidence were repeatedly reinterpreted, reconstructed, or lost entirely at lifecycle transitions. This observation motivated the central research question: *What architectural characteristics enable preservation of manufacturing meaning across the product lifecycle in regulated environments?*

3.2. Objectives of Solution

- Based on problem analysis, the research established four primary objectives for the architectural solution:

2. **Preservation of Manufacturing Intent**: The architecture must support explicit, versioned, and traceable representation of what should be manufactured, under what constraints, and how success is determined—independent of any specific execution system or site.
3. **Semantic Continuity Through Ontology**: The architecture must provide formal mechanisms for defining and governing the meaning of manufacturing concepts (processes, materials, equipment, quality attributes) such that interpretation remains consistent across systems, sites, and lifecycle phases.
4. **Event-Driven Capture of Execution Reality**: The architecture must represent manufacturing execution as time-coherent events and state transitions, binding these events to relevant process, material, equipment, and quality context at the point of generation.
5. **Progressive Evidence Accumulation**: The architecture must structure execution data as quality-relevant evidence with preserved lineage, provenance, and integrity, suitable for continuous disposition rather than retrospective reconstruction.

3.3. Design and Development

- The Intent-Driven Digital Core framework was developed through iterative synthesis of regulatory requirements, international standards, and industry best practices. The design process proceeded through three stages:

- **Stage 1: Regulatory and Standards Analysis**

- Systematic review of regulatory guidance from FDA, ICH, EMA, and WHO established baseline requirements for data integrity, process validation, and lifecycle management. Analysis of international standards including ISA-88 (batch control), ISA-95 (enterprise-control integration), and IEC 62264 provided manufacturing domain abstractions. ICH Q8 (pharmaceutical development), Q9 (quality risk management), Q10 (pharmaceutical quality system), and Q12 (lifecycle management) established quality system expectations.

- **Stage 2: Ontological Foundation Development**

- Through active participation in the Pistoia Alliance CMC Process Ontology project, the research engaged with pharmaceutical manufacturers, technology vendors, and standards bodies to develop shared semantic representations of manufacturing processes. This work established that semantic interoperability—not merely syntactic data exchange—was essential for lifecycle continuity. The ontological approach was further validated through examination of complementary initiatives including NIIMBL Biopharmaceutical

Manufacturing Ontology and NIST Industrial Ontologies Foundry (IOF) work. Figure 4 illustrates this ontology landscape, showing how foundational standards build toward domain-specific applications.

- **Stage 3: Architectural Construct Synthesis**
- The Intent-Driven Digital Core emerged as a unifying construct integrating four interdependent capabilities: semantic definition (ontology-based meaning), event representation (time-coherent execution reality), context binding (process/material/equipment/quality relationships), and evidence structuring (progressive quality-relevant accumulation). The architecture was deliberately positioned as technology-agnostic, defining responsibilities and relationships rather than prescribing specific implementations.

3.4. Demonstration

- The framework's utility was demonstrated through multiple mechanisms:
- **Industry Consortium Validation**: The Pistoia Alliance CMC Process Ontology project (involving AstraZeneca, GSK, Johnson & Johnson, Eli Lilly, Merck, Amgen, and technology partners) provided real-world validation of the semantic layer concepts. The ontology's ability to represent processes across small molecules, biologics, and advanced therapy medicinal products (ATMPs) demonstrated practical applicability.
- **Consensus Document Alignment**: The BioPhorum Digital Batch Disposition Manifesto (co-authored with industry participants including GSK, Vertex, Lonza, Abbvie, and technology vendors) validated the evidence accumulation and continuous disposition concepts. The manifesto's industry endorsement demonstrated consensus on architectural approaches to manufacturing intelligence.
- **Reference Architecture Mapping**: The framework was mapped to established architectural patterns including event-driven architecture, domain-driven design, and data mesh principles, demonstrating conceptual soundness and alignment with broader architectural thinking.

3.5. Evaluation

- Evaluation of the Intent-Driven Digital Core framework occurred through multiple lenses:
- **Regulatory Alignment Assessment**: The architecture was systematically evaluated against regulatory expectations for data integrity (21 CFR Part 11, EU GMP Annex 11, ALCOA+ principles), process validation (FDA 2011 guidance), and lifecycle management (ICH Q12). This assessment confirmed that the architectural approach not only supports compliance but enables more effective demonstration of control and understanding.
- **Standards Conformance Analysis**: Alignment with ISA-88/95 object models, IEC 62264 integration standards, and emerging AI governance frameworks (NIST AI RMF, ISO/IEC 42001) was systematically verified. The architecture was positioned as an extension of—not replacement for—established standards.
- **Industry Consensus Evaluation**: The framework's consistency with ISPE Pharma 4.0 principles, BioPhorum data enablement recommendations, and World Economic Forum manufacturing digitalization research provided external validation of practical relevance.

- **Architectural Stress Testing**: The framework was evaluated against challenging use cases including multi-site technology transfer, CDMO partnerships, continuous manufacturing scenarios, and AI model deployment in regulated contexts. These stress tests revealed where architectural clarity was most essential and where organizational capabilities required explicit design attention.

3.6. Communication

- This paper represents the primary communication artifact, translating industry-grounded work into academic contribution. Additional communication occurred through:
- Consortium workshop presentations (Pistoia Alliance CMC Ontology workshops, 2024-2025)
- Industry working group participation (BioPhorum Data Enablement Program)
- Technical documentation contributions (Batch Disposition Manifesto)
- The framing deliberately bridges industry practice and academic rigor, recognizing that architectural frameworks must be both theoretically sound and practically implementable. This integration reflects the author's direct experience: over a decade working across pharmaceutical and biotechnology organizations on manufacturing systems, process automation, and digital transformation initiatives has informed every aspect of this research—from problem identification through architectural design to evaluation of practical constraints.

4. DATA SOURCES AND EVIDENCE BASE

- The research draws upon multiple categories of sources:

4.1. Primary Regulatory and Standards Sources

- **Regulatory Guidance**: FDA (process validation, data integrity, AI/ML), ICH (Q8, Q9, Q10, Q12, Q5E), EMA (GMP Annex 11), WHO (technology transfer), PIC/S (data management)
- **International Standards**: ISA-88 (IEC 61512), ISA-95 (IEC 62264), IEC 62541 (OPC UA), ASTM E2500 (pharmaceutical systems)
- **Ontology Standards**: ISO/IEC 21838-2 (Basic Formal Ontology), IOF Core Ontology, QUDT (units and quantities)

4.2. Industry Consortium Materials

- **Pistoia Alliance**: CMC Process Ontology project documentation, workshop presentations, phase reports, ontology release artifacts
- **BioPhorum**: Data Enablement for AI reports, "What Is AI-Ready Data?" guidance, Digital Plant Maturity Model, Digital Batch Disposition Manifesto
- **ISPE**: Pharma 4.0 Operating Model, GAMP 5 (2nd Edition), Pharma 4.0 Digital Maturity Model, AI in Pharmaceutical Manufacturing guidance
- **NIIMBL**: Biopharmaceutical Manufacturing Ontology, big data program outputs
- **NIST**: Industrial Ontologies Foundry materials, semantic interoperability research, AI Risk Management Framework

4.3. Scholarly Literature

- Design science research methodology (Hevner, March, Peffers, Gregor)
- Enterprise architecture frameworks (Zachman, TOGAF)

- Event-driven architecture patterns (Hohpe & Woolf, Fowler)
- Ontology engineering (Gruber, Guarino, Arp et al.)
- AI governance and lifecycle management (Mitchell, Gebru, Sculley)
- Manufacturing systems and quality (ISA, IEC, pharmaceutical engineering literature)

4.4. Industry Research, Analysis, and Practical Manufacturing Experience

- McKinsey & Company, Deloitte, PwC: Digital transformation studies, pharmaceutical operations research
- World Economic Forum: Global Lighthouse Network reports, Industry 4.0 manufacturing assessments
- Gartner, Forrester: Technology landscape analysis, digital maturity frameworks
- Practical manufacturing experience: Direct involvement in manufacturing systems implementation, process automation, batch record design, technology transfer activities, and digital transformation initiatives across pharmaceutical and biotechnology organizations over more than a decade

5. SYNTHESIS APPROACH

- The research employed a structured synthesis methodology integrating regulatory requirements, industry standards, consortium insights, and scholarly literature into a coherent architectural framework.

- **Stage 1: Requirements Extraction**

- Regulatory guidance and standards were systematically analyzed to extract explicit and implicit architectural requirements. For example, FDA data integrity guidance (2018) implies requirements for attribution, contemporaneous capture, and immutable recording —which translate to architectural responsibilities for provenance tracking, event timestamp binding, and append-only evidence structures.

- **Stage 2: Gap and Discontinuity Analysis**

- Industry consortium participation and manufacturing documentation review revealed recurring patterns of semantic loss, contextual collapse, and evidence reconstruction. These patterns were mapped to specific lifecycle transition points (development to manufacturing, site to site, internal to CDMO) and categorized by architectural cause (missing ontology, weak event coherence, inadequate context binding).

- **Stage 3: Architectural Pattern Integration**

- Established architectural patterns (event sourcing, domain-driven design, ontology-based data access) were evaluated for applicability to pharmaceutical manufacturing constraints. The Intent-Driven Digital Core emerged as a domain-specific instantiation of broader architectural thinking, adapted for regulated lifecycle environments.

- **Stage 4: Iterative Refinement**

- The framework was iteratively refined through consortium engagement, particularly testing ontological approaches against multi-modality manufacturing requirements and validating evidence structures against batch disposition workflows.

6. CONTRIBUTION TO KNOWLEDGE

- This research makes three primary contributions:

1. **Architectural Framework**: The Intent-Driven Digital Core provides a novel reference architecture specifically addressing lifecycle continuity challenges in regulated pharmaceutical manufacturing, bridging gaps between existing standards (ISA-88/95), regulatory expectations, and emerging digital capabilities.
2. **Ontology as Architectural Mechanism**: The research positions ontology not as an academic exercise but as a practical architectural construct essential for semantic interoperability across systems and lifecycle phases, validated through industry consortium development.
3. **Socio-Technical Integration**: By explicitly treating people, roles, and capabilities as architectural components rather than external users, the framework provides a more complete model of digital manufacturing systems in regulated environments.

7. REFERENCES FOR METHODOLOGY SECTION

- Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, 37(2), 337-355.
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75-105.
- March, S. T., & Smith, G. F. (1995). Design and natural science research on information technology. *Decision Support Systems*, 15(4), 251-266.
- Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45-77.
- Simon, H. A. (1996). *The Sciences of the Artificial* (3rd ed.). MIT Press.
- van Aken, J. E. (2004). Management research based on the paradigm of the design sciences: The quest for field-tested and grounded technological rules. *Journal of Management Studies*, 41(2), 219-246.

-
-
-