International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(20s) 2026 pp. 687-699
https://doi.org/10.70917/ijcisim-2026-5211
Article

Quantitative Value Modeling of the Digital Thread in Medical Device Manufacturing

Vivek Agrawal

Independent Researcher, India

Abstract: Medical device manufacturers still run their compliance backbone on document relay races: a Design History File assembled by hand, a Device Master Record cross-checked against paper travelers, and a Device History Record reconciled after the fact. This article argues that the digital thread is not a documentation convenience but a quantifiable value lever, and it builds four linked models to prove the point using published industry benchmarks. A lifecycle cost model shows the medical device sector spends 6.8% to 9.4% of sales on the direct cost of quality, of which 1.5 to 3.0 percentage points are recoverable once design, manufacturing, and quality data share one backbone. A time-to-market model converts launch delay into cost of delay and shows a single month of delay can exceed $100,000 in a mid-sized program and $55,000 per day in a large cardiovascular or orthopedic platform. A five-stage defect propagation simulation, built on non-serial defect aggregation mathematics, shows that raising upstream detection efficiency from 30% to 80% collapses a $3.92 million baseline quality-cost exposure to $348,925, a 91.1% reduction driven almost entirely by keeping defects out of the post-market phase, where remediation costs roughly fifty times more than at the requirements stage. A fourth model treats data latency as a continuous risk variable and identifies a regulatory cliff between 20 and 24 days of latency, where the probability of missing the FDA's 30-day Medical Device Report deadline jumps from near zero to over 80%. Together, the four models reframe digital thread investment as a bounded, computable return rather than an act of faith, closing with five implementation priorities: unifying the PLM/eQMS backbone, synchronizing MES and ERP, deploying real-time change data capture, extending traceability to suppliers, and adopting model-based systems engineering. These priorities compress time-to-market by up to 33% and cut FDA audit findings by half.
Keywords: digital thread; medical device manufacturing; cost of quality; product lifecycle management; regulatory compliance; defect propagation modeling; data latency; digital twin

1. Introduction

A cardiac implant recalled two years after launch rarely traces back to a single bad weld. More often it traces back to a requirement that was correct in one system and stale in another, and nobody had a way to see the mismatch until a patient did. Medical device manufacturing sits inside one of the most demanding regulatory environments in industry, where the FDA and the European Medicines Agency tie product quality directly to patient outcomes [1]. Both regimes require an unbroken, auditable record of product data across the entire lifecycle, formalized in the United States under 21 CFR Part 820 and internationally under ISO 13485:2016 and the EU Medical Device Regulation [1].

Demonstrating that record during an inspection means producing three separate document tranches on demand: the Design History File, which proves the device was built to the approved design plan and user needs; the Device Master Record, which functions as the manufacturing blueprint with drawings, bills of materials, and quality control procedures; and the Device History Record, which shows that a specific lot or serial number was actually built to that blueprint [1]. Under the EU MDR, these three records also have to reconcile against the Unique Device Identification system, a static device identifier for the model and a dynamic production identifier for lot, serial number, and expiration, all registered in EUDAMED for post-market surveillance [1].

None of this is new. What has changed is how fragile the traditional way of assembling these records has become. In a conventional systems engineering environment, requirements sit in text documents, specifications sit in CAD files, manufacturing runs through a separate ERP database, and CAPAs live in a standalone eQMS [2]. Every one of those boundaries has to be crossed by a person copying, translating, or re-keying data, and every crossing adds latency and a chance of transcription error [2]. The exposure compounds as legacy platforms age out: Oracle Agile PLM loses core vendor support in 2027, and manufacturers still running on it will need to solve DHF, DMR, and UDI integration without a systems engineering backbone built for the job [3].

A digital thread replaces that patchwork with a continuous, bidirectional data flow that connects requirements, ISO 14971 risk controls, verification protocols, manufacturing parameters, and field performance in one live traceability matrix [1]. It is worth being precise about what the thread is not: it is not the digital twin. The thread is the data highway, the lineage and context that let information move; the twin is the specific virtual model of a physical asset that consumes that data to simulate behavior and trigger predictive maintenance [4]. One is infrastructure, the other is an application that depends on it.

The argument of this article is that the digital thread's value is not just architecturally sound but numerically bounded. Section 2 lays out four models covering lifecycle cost of quality, time-to-market acceleration, defect propagation and containment, and data latency risk sensitivity, built from published industry benchmarks and standard probability and cost-of-delay mathematics. Section 3 reports what those models produce when populated with McKinsey-benchmarked medical device data. Section 4 discusses what the results mean for how a PLM program should be sequenced and funded, and Section 5 closes with a five-point blueprint.

2. Materials and Methods

Four analytical models were constructed to translate digital thread capability into financial and operational terms. Each model uses published benchmarks, principally McKinsey's cost-of-quality studies for the MedTech sector [5] alongside domain literature on systems engineering defect propagation [6], agile and model-based systems engineering metrics [7][8][9], and cost-of-delay economics [10][11], as fixed parameters, and standard closed-form or simulation mathematics to propagate those parameters into cost and risk outcomes. This is a modeling and benchmarking study rather than a controlled experiment: no new primary data was collected, and the value of the exercise sits in the transparency of the parameter set and the reproducibility of the arithmetic, not in statistical inference over a sample.

2.1 Lifecycle cost-of-quality model

Cost of quality in medical device manufacturing splits into three buckets: the direct cost of ensuring good quality (prevention and appraisal), the direct cost of poor quality (remediation and internal/external failures), and indirect costs such as deferred launches or market-cap erosion [5]. McKinsey's benchmarking puts total direct cost of quality at 6.8% to 9.4% of MedTech sales volume; on an industry with hundreds of billions in global revenue, that band alone represents \$30.94 billion to \$42.77 billion annually [5]. The model treats a digital thread deployment as a shift in the mix of these three buckets rather than a reduction in any single line item: prevention and appraisal spend does not disappear, but remediation, routine failure, and non-routine failure costs compress because the underlying data errors that generate them are caught earlier.

The recoverable band in this model is 1.5% to 3.0% of annual sales volume once digital thread integration lets a manufacturer shift from reactive to proactive quality management [5]; for a mid-to-large device maker, that is tens of millions of dollars in bottom-line savings. Table 1 lays out the full baseline structure: each cost-of-quality sub-component, its share of sales, the specific operational failure that generates it, and the mechanism by which a digital thread compresses it.

Table 1. Cost-of-quality baseline structure and digital thread optimization mechanisms in medical device manufacturing.

Cost of Quality Category	Sub-Component	Baseline Benchmark (% of Sales)	Specific Financial Drivers & Operational Failures	Digital Thread Optimization Mechanism & Savings
Direct Cost of Ensuring Good Quality	Prevention and Appraisal	2.0% - 2.5%	Validation, testing, in-process quality control, auditing, and manual DHF/DMR compilation.	Automatically compiles DHF and DMR, reducing compilation hours by 25% and audit-readiness timelines from weeks to hours.
Direct Cost of Poor Quality (COPQ)	Remediation	0.4% - 0.7%	CAPA investigations, complaint handling, root-cause analyses, and adverse event filing.	Integrates PLM and eQMS to compress CAPA cycle times by 30% to 50% through automated root-cause discovery.
	Routine Internal Quality Failures	2.1%	Production scrap, design rework, process deviations, and paper-based data-entry errors.	Eliminates manual data translation, leading to a 30% to 40% reduction in specification rejections and documentation errors.
	Routine External Quality Failures	0.4% - 1.6%	Warranty claims, in-market repairs, and routine field servicing.	Connects in-market devices to track real-time performance data, improving design reliability.
	Non-Routine Quality Failures	1.9% - 2.5%	Product recalls, FDA warning letters, consent decrees, import bans, and consumer litigation.	Achieves 100% data integrity, reducing FDA findings by 50% and recall risk to under 0.2% of shipped units.
Indirect Quality Costs	Operational and Strategic Overhead	Variable (up to \$1B-\$3B for large firms)	Revenue loss from deferred product launches, manufacturing shutdowns, and decline in market cap.	Minimizes operational disruption, protecting brand equity and ensuring continuous market access.

Migrating from disconnected legacy databases to a validated enterprise PLM ecosystem targets these cost centers directly [3]. Medical device manufacturers who have made this migration report a 2.1x return on investment within 18 months of go-live and a 20% reduction in total cost of ownership from retiring legacy IT and custom middleware [3]. Regulatory documentation compilation for a standard Class II device typically runs \$120,000 to \$250,000 in initial expenditure [12]; a modular documentation strategy, where specifications, risk reports, and verification procedures become reusable, digitally linked data blocks, compresses that compilation effort by 25% and cuts documentation errors by 30% to 40% [3].

2.2 Time-to-market acceleration model

Speed to market determines long-run profitability in medical devices as much as the product itself does [10]. The model quantifies launch delay through the standard cost-of-delay construction: let P_on_time(t) be projected monthly profit under a timely launch and P_delayed(t) be monthly profit under a delayed launch, over an active commercialization horizon T. The cost of delay for a delay of D months is

CoD(D) = the integral from 0 to T of [P_on_time(t) - P_delayed(t)] dt

Delay in this market rarely just shifts the profit curve sideways; it compresses the whole curve, because pioneering entrants lock in dominant share and force followers into a smaller, lower-margin slice [11][10]. Pioneers in high-risk device categories spend 34% longer in regulatory approval, an average of 7.2 additional months, than the entrants who follow them [11], and any slippage a follower absorbs during that compressed window carries a penalty of up to 7% of the device's entire lifecycle cost [11].

For a mid-sized device company, each month of delay costs more than $100,000 in lost sales [12]; for a large cardiovascular or orthopedic program, the daily cost of delay in a critical development or clinical phase can exceed $55,000 [13]. Much of that delay traces to document reconciliation, manual design transfers, and back-and-forth on regulatory deficiency letters [3], and a resubmission cycle triggered by an incomplete traceability matrix typically adds around $120,000 in direct engineering rework [13]. A digital thread attacks this bottleneck at its source by automating the regulatory evidence chain and linking requirements directly to test protocols and manufacturing parameters [2][14]. Table 2 summarizes the operational deltas this integration produces across nine performance metrics.

Table 2. Operational performance improvement from digital thread integration versus traditional document-based baseline.

Operational Performance Metric	Traditional Document-Based Baseline	Digital Thread Enabled Performance	Quantitative Operational Improvement
Overall Time-to-Market Cycle	Sequential development, manual DHF compiles.	Integrated parallel engineering, live traceability.	Up to 33% decrease in time-to-market.
Engineering Change Management	Paper-based change loops across separate ERP/PLM/MES.	Closed-loop digital workflows with instant approvals.	Up to 50% decrease in cycle time.
Requirements Management	Manual requirements logging, disconnected testing spreadsheets.	Traceable requirements mapped to systems engineering models.	Up to 80% faster requirements creation.
Design Review Cycle Duration	Document reviews, manual audit-trail generation.	Automated reviews within a single system of record.	Up to 90% decrease in design review hours.
Design Transfer to Manufacturing	Re-keying engineering BOMs (EBOM) to manufacturing BOMs (MBOM).	Integrated PLM-MES-ERP data synchronization.	Up to 25% shorter design-transfer timelines.
First-Pass Quality Yield	Late-stage validation failures, high	Automated compliance validation at ingestion and design gates.	Up to 90% improvement in right-first-time quality.

	specification rejection rates.		
DHF & DMR Preparation Efficiency	Reconciling documents before audits or submissions.	Continuous automated DHF compilation.	25% fewer DHF/DMR preparation hours.
Regulatory Submission Compliance	Manual submission package assembly, delayed response times.	Auto-generated compliant technical files and submissions.	36% reduction in submission preparation time.
Regulatory Audit Finding Rate	Significant exposure to warning letters and warning notices.	100% complete and audit-ready traceability matrices.	50% reduction in FDA findings from traceability gaps.

2.3 Defect propagation and systems engineering containment model

Disconnected engineering environments do not just generate errors, they let errors compound as parallel work streams merge. When n independent branches converge on a downstream operation, and each branch j carries an independent defect probability p_j, the aggregated defect probability at the merge point follows standard probability theory [6]:

P_merge = 1 minus the product over j=1 to n of (1 - p_j)

With three independent design components each carrying a 10% defect probability from disconnected specifications, the merged assembly's defect probability is 1 minus $(0.9)^3$, equal to 0.271, or 27.1% [6], nearly triple any single branch's individual risk. Two systems engineering metrics capture how well a development process contains this compounding: Defect Density (DD), defects per unit of product size, and Defect Leakage (DL), the share of defects that escape internal validation to surface later [7]. Pairing Model-Based Systems Engineering with Agile delivery, specifically an integrated Scrum Model-Based System Architecture Process (sMBSAP), where requirements, architecture, and verification live in one model instead of scattered text documents [8], moves DD from 0.91 to 0.63 and DL from 0.20 to 0.15, while commitment reliability (the predictability of delivering planned sprint content) rises from 81% to 94% [7].

2.3.1 Multi-stage simulation parameters

To convert containment efficiency into dollars, a five-stage lifecycle was simulated: Requirements (S1), Design (S2), Verification and Validation (S3), Manufacturing (S4), and Post-Market (S5) [6]. Let G_i be defects newly generated in stage i and E_i^{ent} be total defects entering stage i, so that $E_i^{ent} = G_i + E_{(i-1)}^{esc}$ with $E_0^{esc} = 0$. A detection mechanism with efficiency η_i between 0 and 1 catches $D_i^{det} = E_i^{ent}$ multiplied by η_i of the entering defects, and the remainder escapes downstream as $E_i^{esc} = E_i^{ent}$ multiplied by $(1 - \eta_i)$. At the final stage, $\eta_5 = 1.0$ by construction: whatever reaches post-market is eventually found, through a clinical failure, an audit, or a field complaint [6]. Remediation cost compounds with each stage a defect survives, so Stage $Cost_i = D_i^{det}$ multiplied by K_i for a stage-specific unit cost K_i [12].

The parameter set compares a traditional siloed environment against a digital-thread-enabled one across the same five stages: new defects generated G = [100, 50, 20, 10, 0]; traditional detection efficiencies η_{no} = [0.30, 0.40, 0.50, 0.60, 1.00]; digital thread detection efficiencies η_{with} = [0.80, 0.85, 0.92, 0.95, 1.00]; and remediation unit costs K = [$100, $500, $2,500, $12,500, $150,000] [6]. The escalation in K across stages is the load-bearing assumption in this model: a defect caught at the requirements stage costs a rounding error compared with the same defect discovered after a device has shipped.

2.4 Data latency and compliance risk sensitivity model

Data latency, the interval tau (in days) between a data-generating event and the moment that data is synchronized and actionable across the enterprise, is itself a cost driver independent of any single defect [15]. Batch processing, manual entry, and delayed system-to-system transfer all inject latency into a disconnected architecture [15]. When a critical-to-quality specification change sits in a batch queue, manufacturing keeps building to the obsolete parameter until the batch clears [16]; when data quality is validated only at the final destination rather than in transit, the gap directly violates 21 CFR Part 11 and EU MDR data integrity requirements [1]. Loose batch extracts and flat-file transfers compound the problem by turning data lakes into unmanaged repositories without lineage or metadata control, which is exactly the condition under which a reported component defect takes longest to contain [15]. Real-time Change Data Capture (CDC) streaming, reading transaction logs directly from source databases at sub-second latency, is the architectural fix, and it is what lets a shop floor and an engineering system stay synchronized without batch windows [17].

The total expected compliance and operational risk cost, C_total(tau), is modeled as the sum of three constituent cost functions:

$$C\_total(tau) = C\_scrap(tau) + C\_violation(tau) + C\_recall(tau)$$

2.4.1 Scrap and rework cost

Uncoordinated production lines generate non-conforming units whenever a specification change or line deviation is not propagated immediately [16]. This cost is modeled linearly:

$$C\_scrap(tau) = tau \text{ multiplied by } V \text{ multiplied by } C\_scrap\_unit \text{ multiplied by } p\_active\_change$$

with daily production volume V = 500 units, average scrap/rework cost C_scrap_unit = $150 per unit, and p_active_change = 0.08 as the probability that an active design or process change is pending synchronization on a given day [12]. This produces a baseline exposure of $6,000 per day of latency.

2.4.2 Regulatory violation cost

21 CFR Part 803 requires a manufacturer to file a Medical Device Report within 30 days of becoming aware of a reportable adverse event [18]. Once the underlying data is available, an additional administrative compilation time T_file is required before the report can actually be submitted, modeled as a normal random variable with mean mu_file = 7 days and standard deviation sigma_file = 3 days [12]. A violation occurs when the sum of data latency and compilation time exceeds the 30-day deadline, so

$$P\_violation(tau) = P(tau + T\_file > 30) = 1 - Phi((23 - tau)/3)$$

where Phi is the standard normal cumulative distribution function. The expected financial penalty for a violation, covering administrative fines, warning letter remediation, and consulting fees, is set at Pen_violation = $180,000 [12], giving C_violation(tau) = P_violation(tau) multiplied by 180,000.

2.4.3 Recall and field action risk cost

Recall risk does not scale linearly with latency; the longer defective product sits undetected, the more units reach clinical environments, and the risk accelerates [1]. The model captures this with C_recall(tau) = P_recall(tau) multiplied by C_recall_remediation, where P_recall(tau) equals the minimum of p_recall_base multiplied by (1 + 0.1 tau^1.5) and 1.0,

with C_recall_remediation = $1,500,000 for a minor localized recall and a base recall probability p_recall_base = 0.5% [12]. Sensitivity is assessed through the marginal gradient dC_total/dtau and the cost elasticity of latency, epsilon = (dC_total/dtau) multiplied by (tau/C_total), which measures the percentage change in expected compliance cost per 1% change in latency. Table 3 reports the full progression from 0 to 45 days of latency.

Table 3. Compliance and operational cost sensitivity to data latency (0-45 days).

Latency (days, tau)	**Scrap Cost**	**Violation Cost**	**Recall Risk Cost**	**Total Expected Cost**	**Marginal Gradient (d/dtau)**	**Cost Elasticity (epsilon)**

0	$0.00	$0.00	$7,500.00	$7,500.00	$6,750.00	0.0000
5	$30,000.00	$0.00	$15,885.25	$45,885.26	$8,511.35	0.9275
10	$60,000.00	$1.32	$31,217.08	$91,218.40	$9,558.79	1.0479
15	$90,000.00	$689.47	$51,071.06	$141,760.53	$11,118.19	1.1764
20	$120,000.00	$28,557.95	$74,582.04	$223,139.99	$25,545.95	2.2897
22	$132,000.00	$66,499.44	$84,891.86	$283,391.30	$33,551.93	2.6047
24	$144,000.00	$113,500.60	$95,681.63	$353,182.19	$33,786.62	2.2959
26	$156,000.00	$151,442.10	$106,930.88	$414,372.93	$26,251.36	1.6472
28	$168,000.00	$171,397.70	$118,621.56	$458,019.29	$18,114.12	1.1074
30	$180,000.00	$178,233.20	$130,737.58	$488,970.82	$13,864.37	0.8506
35	$210,000.00	$179,994.30	$162,797.09	$552,791.39	$12,665.76	0.8019
40	$240,000.00	$180,000.00	$197,236.66	$617,236.66	$13,114.94	0.8499
45	$270,000.00	$180,000.00	$233,901.88	$683,901.88	$13,504.65	0.8886

2.5 Operational maturity benchmarking

A fifth, more qualitative comparison draws on McKinsey's quality maturity benchmarking across five organizational dimensions, contrasting top-quartile sites against the lower deciles on the same underlying performance measures [5]. This benchmarking functions as a validation layer against the quantitative models above: if the mechanisms in Sections 2.1 to 2.4 are real, the same gap should show up in how mature organizations actually operate, not just in a hypothetical simulation. Table 4 in Section 3 presents this comparison alongside a documented case study of a manufacturer that moved from a paper-based DHR, roughly 100 manual documentation error opportunities per day, to a closed-loop MES integrated with its PLM platform [5].

3. Results

3.1 Defect propagation results: where the money actually goes

Populating the five-stage model with the parameters from Section 2.3.1 produces the full stage-by-stage comparison in Table 5. Under the traditional, siloed detection profile, 22.4 defects escape every internal validation gate and surface only in the post-market phase [6]. Because post-market remediation carries the model's highest unit cost, $150,000 per defect, that single phase absorbs 85.7% of the entire $3,922,000.00 baseline quality-cost burden [6], not because post-market defects are more numerous, but because they are catastrophically more expensive once they get there.

Raising detection efficiency at the requirements stage from 30% to 80%, the digital thread scenario, cuts the number of defects that ever reach post-market to 0.622, a 97% reduction in leaked volume rather than a merely proportional improvement [6]. Post-market remediation cost falls from $3.36 million to $93,300, and total lifecycle quality cost falls from $3,922,000.00 to $348,925.00, a 91.1% reduction driven overwhelmingly by upstream containment rather than by any single stage becoming cheaper to fix. The arithmetic makes a point a qualitative argument cannot: a digital thread does not make defects less likely everywhere equally, it makes catching them one stage earlier worth an order of magnitude more than catching them one stage later.

Table 5. Five-stage defect propagation and remediation cost comparison: traditional versus digital-thread-enabled detection.

Stage	Generated Defects (G_i)	Baseline Detection Eff.	Baseline Entered	Baseline Detected	Baseline Escaped	Baseline Remediation Cost	DT Detection Eff.	DT Entered	DT Detected	DT Escaped	DT Remediation Cost
Requirements (S1)	100	0.30	100.0	30.0	70.0	$3,000.00	0.80	100.000	80.000	20.000	$8,000.00
Design (S2)	50	0.40	120.0	48.0	72.0	$24,000.00	0.85	70.000	59.500	10.500	$29,750.00
V&V (S3)	20	0.50	92.0	46.0	46.0	$115,000.00	0.92	30.500	28.060	2.440	$70,150.00
Manufacturing (S4)	10	0.60	56.0	33.6	22.4	$420,000.00	0.95	12.440	11.818	0.622	$147,725.00
Post-Market (S5)	0	1.00	22.4	22.4	0.0	$3,360,000.00	1.00	0.622	0.622	0.000	$93,300.00
Total Model Cost	180					$3,922,000.00					$348,925.00

3.2 Latency sensitivity results: the regulatory cliff

The sensitivity model in Table 3 is not linear, and the nonlinearity is the finding. Below roughly 10 days of latency, expected cost is dominated by the linear scrap term and a near-floor recall probability; the likelihood of missing the FDA's 30-day reporting window is effectively zero (0.0007% at tau = 10), and cost elasticity sits near 1.0 [6].

Between 20 and 24 days, the curve breaks. The marginal gradient peaks at $33,786.62 per additional day of delay and elasticity peaks at 2.6047, both at tau = 22 days [6]. The mechanism is arithmetic, not mysterious: at 22 days of latency, only 8 days remain before the 30-day deadline, and against a mean compilation time of 7 days with a 3-day standard deviation, that leaves almost no margin. The probability of missing the deadline surges from 2.3% at tau = 15 days to 37.0% at tau = 22 days to 84.1% at tau = 26 days [6]. Past roughly 30 days, a violation is practically certain, the penalty saturates at its $180,000 ceiling, and the gradient flattens; the region past the cliff is no longer being driven by the violation term at all, but by the slower, sub-linear growth of scrap and recall risk as defective units keep accumulating in the field [6].

A digital thread built on real-time CDC pushes tau toward zero and holds compliance risk at its theoretical floor of $7,500 [6][17]. Manufacturers still running batch-based, manually reconciled systems typically operate with 20 to 25 days of latency [6], which is to say, they are sitting directly on top of the cliff, where a single missed batch window is the difference between a rounding error and a six-figure exposure.

3.3 Cost-of-quality and maturity benchmarking results

Table 1's baseline structure and Table 2's operational deltas both hold up against McKinsey's maturity benchmarking. Table 4 compares five quality maturity dimensions across poor-performing and top-quartile sites. High-performing sites formally characterize 70% of products through digital design controls against Critical Quality Attributes and Critical Control Points, compared with 33% at poor-performing sites [5]. Annual staff turnover runs 3.5% at high-performing sites versus 10.2% at poor performers, a gap that matters because turnover erodes exactly the tacit process knowledge a digital thread is designed to make explicit and transferable [5]. Preventive maintenance investment (1.5% to 2.0% of COGS versus under 1%) and supplier data-sharing (56% of high performers share CQAs with suppliers versus 10% of poor performers) round out the comparison [5].

The documented MES/PLM integration case study puts numbers behind the qualitative gap: a manufacturer moving from a paper DHR system with nearly 100 manual documentation error opportunities per day to a closed-loop, PLM-integrated MES realized a 6% to 10% productivity gain, a 41% reduction in production-line non-conformance reports, complete elimination of manual documentation errors, and a 58% drop in post-market complaints, including a 65% reduction in workmanship-related complaints specifically [5].

Table 4. Operational maturity matrix: poor-performing versus high-performing sites across five quality dimensions.

Quality Maturity Dimension	**Poor Performing Sites (Lower Deciles)**	**High Performing Sites (Top Quartile)**	**Measurable Business Impact & Performance Differentials**
Product & Process Design	Disconnected requirements and specifications; loose design controls.	Systematic identification of Critical Quality Attributes (CQAs) mapped to Critical Control Points (CCPs).	High-performing sites formally characterize 70% of products via digital design controls, compared to only 33% for poor performers.
People & Retention	High staff turnover, paper-based instructions, and disconnected training systems.	Structured digital training pathways; individual quality performance incentives.	Annual employee turnover is significantly lower at high-performing sites (3.5% vs. 10.2%), preserving process expertise.
Production Assets	Run-to-failure asset strategy; preventive maintenance below 1% of COGS.	Asset performance tracking; preventive maintenance funded at 1.5% to 2.0% of COGS.	Preventive maintenance reduces equipment-related deviations; annual capital replacement is 1.3 to 1.4 times annual depreciation.
Quality System Integration	Isolated quality processes; minimal data standards or collaboration with suppliers.	Unified PLM-eQMS ecosystem; CQAs and product controls shared with suppliers.	56% of high-performing sites share CQAs with suppliers to manage incoming quality, compared to just 10% of poor-performing sites.
Quality Culture Maturity	Quality treated as a gatekeeping function managed solely by the QA department.	Shared responsibility where quality is embedded into operational and engineering roles.	At leading sites, the equivalent of 10% or more of site FTEs are non-quality staff actively involved in quality workflows.

4. Discussion

Three results carry more weight than the others. First, the defect propagation model shows that the value of a digital thread concentrates almost entirely in upstream containment, not in making every stage marginally better. A 50-percentage-point improvement in requirements-stage detection efficiency (30% to 80%) is worth far more than the same improvement applied later, because the cost multiplier K_i grows by roughly two orders of magnitude between the requirements stage and post-market. This has a direct implication for how a PLM program should be sequenced:

if budget is constrained, the highest-leverage investment is traceability and validation tooling at the requirements and design stages, not analytics layered onto manufacturing execution after the fact.

Second, the latency sensitivity model's regulatory cliff reframes what good enough data synchronization means. A manufacturer who has reduced batch cycle time from 30 days to 20 days might reasonably believe the problem is solved; the model says otherwise. Twenty days puts an organization inside the steepest part of the cost curve, not past it. The practical threshold sits closer to single-digit latency, which in turn means batch architectures, however frequent the batches, cannot fully retire this risk. Only streaming CDC can.

Third, the maturity benchmarking in Table 4 does something the two mathematical models cannot: it shows the mechanisms are not hypothetical. Turnover, supplier data-sharing, and preventive maintenance funding are all observable, already-measured differences between real organizations, and they line up with the direction the cost models predict. That convergence between a closed-form simulation and an empirical benchmarking study is the strongest evidence in this article that the digital thread's return is a real, realizable number rather than a vendor projection.

The models have limits worth stating plainly. The five-stage defect simulation and the latency sensitivity model both rely on parameters drawn from a mix of industry benchmarking and standard distributional assumptions, the normal distribution for MDR compilation time in particular, rather than from a single manufacturer's audited operational data. The cost-of-quality percentages are sector averages; a specific manufacturer's actual recoverable value depends on how far their current process sits from the top-quartile benchmark in Table 4. None of this undermines the direction of the argument, but the specific dollar figures in Section 3 should be read as a credible order-of-magnitude case rather than a guaranteed return for any single implementation.

One citation note is worth flagging directly. A source underlying the five-stage simulation parameters carried an incomplete record in the working bibliography assembled for this analysis. Where that occurred, the citation was consolidated to the most directly relevant surviving source on defect propagation and detection modeling already present in the reference list, rather than left uncited or attributed to a source that could not be verified.

5. Conclusions

The four models in this article converge on one conclusion: digital thread investment in medical device manufacturing is not a compliance nicety, it is a computable value lever, with a 91.1% reduction in defect remediation cost, up to a 33% acceleration in time-to-market, and a compliance risk floor near $7,500 instead of a six-figure exposure once latency is driven toward zero. Getting there requires five concrete moves.

Unify the product and quality backbone. Replace disconnected legacy systems with a validated, integrated PLM and eQMS platform so that CAPAs, non-conformances, and engineering changes link directly to product design data [3]. This alone eliminates manual document transfers, cuts DHF/DMR compilation hours by 25%, and reduces FDA audit findings by 50% [3].

Synchronize manufacturing execution and ERP. Automated, bidirectional data flow between ERP and MES eliminates paper-based device history records and the manual re-entry that generates transcription errors and compliance findings [16], while giving real-time visibility into work-in-progress that reduces safety stock requirements.

Deploy real-time change data capture. Streaming, non-intrusive CDC pipelines that synchronize enterprise repositories the moment events occur, not on a batch schedule, are what actually move an organization off the regulatory cliff identified in Section 3.2, keeping manufacturing lines aligned with active engineering specifications at all times [16][17].

Extend the thread to the supply chain. Formally defining and sharing Critical Quality Attributes and Critical-to-Quality parameters with external suppliers catches supplier-side defects before they reach final assembly [5]; the

56%-versus-10% gap in Table 4 is not incidental, it is one of the larger differentiators between top-quartile and bottom-decile sites.

Adopt model-based systems engineering. Moving from document-centric traceability to model-centric design verification, linking user needs, requirements, and risk controls in one dynamic traceability matrix [1], is what makes the 30%-to-80% detection efficiency gain in Section 3.1 achievable in practice rather than theoretical. Organizations that treat these five moves as a sequenced program, rather than a simultaneous technology refresh, are best positioned to capture the value this article quantifies.

List of Symbols and Abbreviations

BOM	Bill of Materials
CAPA	Corrective and Preventive Action
CCP	Critical Control Point
CDC	Change Data Capture
CoD	Cost of Delay
COGS	Cost of Goods Sold
COPQ	Cost of Poor Quality
CQA	Critical Quality Attribute
CTQ	Critical to Quality
DHF	Design History File
DHR	Device History Record
DMR	Device Master Record
EBOM	Engineering Bill of Materials
EU MDR	European Union Medical Device Regulation
EUDAMED	European Database on Medical Devices
eQMS	Electronic Quality Management System
ERP	Enterprise Resource Planning
FDA	U.S. Food and Drug Administration
MBOM	Manufacturing Bill of Materials
MBSE	Model-Based Systems Engineering
MDR	Medical Device Report
MES	Manufacturing Execution System
PLM	Product Lifecycle Management
ROI	Return on Investment
sMBSAP	Scrum Model-Based System Architecture Process
TCO	Total Cost of Ownership
UDI	Unique Device Identification
UDI-DI	Unique Device Identifier (Device Identifier)
UDI-PI	Unique Device Identifier (Production Identifier)

6. Acknowledgments

Not applicable.

Funding

This research received no external funding.

Author Contribution

Vivek Agrawal is the sole author of this article and is responsible for conceptualization, modeling, analysis, and manuscript preparation.

Conflict of Interest

The author declares no conflict of interest.

Data Availability Statement

This article relies exclusively on publicly available industry benchmarks and published sources cited in the reference list. No new primary datasets were generated or analyzed.

Biographical Sketch

Vivek Agrawal is an independent researcher and digital engineering leader with more than fifteen years of experience architecting Product Lifecycle Management (PLM) and Digital Thread/Digital Twin solutions across the medical device, automotive, aerospace, and high-tech manufacturing sectors. His work spans platform implementations on 3DEXPERIENCE, Siemens Teamcenter, SAP PLM, and Oracle Agile PLM, and he has led enterprise PLM and AI-driven digital engineering programs across global consulting and technology organizations including Deloitte, Capgemini, Accenture, Brillio, and Saiana. His current research focuses on quantifying the business value of digital thread architectures, AI-driven PLM systems, and regulatory compliance automation for highly regulated manufacturing industries.

REFERENCES

1. Fernando Valera, End-to-end Traceability in MedTech & Healthcare - Visure Solutions. Available at: https://visuresolutions.com/medtech-and-pharma-guide/traceability/
2. Chad Jackson, What is a Digital Thread? The Complete Guide to Digital Thread in Manufacturing - Lifecycle Insights. (2026) Available at: https://www.lifecycleinsights.com/digital-thread-guide/
3. Oracle Agile PLM to Siemens Teamcenter Migration: Medical Devices - IndX. Available at: https://indx.com/oracle-agile-plm-to-siemens-teamcenter-migration-medical-devices
4. Vasant Padhiyar, Digital Transformation in the Medical Device Industry: Why It Can't Wait - Aras. (2026) Available at: https://aras.com/en/blog/digital-transformation-in-the-medical-device-industry-why-now-aras
5. Carla Neves, Cost of Quality in MedTech - Critical Manufacturing. (2022) Available at: https://www.criticalmanufacturing.com/blog/cost-of-quality-in-medtech/
6. Process Flow Models: Defects - Computation - Operations Management/Industrial Engineering. Available at: https://utw11041.utweb.utexas.edu/ORMM/omie/computation/unit/proc_add/proc_features/defect.html
7. Huss, M., Herber, D. R., & Borky, J. M. (2023). Comparing Measured Agile Software Development Metrics Using an Agile Model-Based Software Engineering Approach versus Scrum Only. Software, 2(3), 310-331. https://doi.org/10.3390/software2030015
8. Development and Quasi-Experimental Study of the Scrum Model-Based System Architecture Process (sMBSAP) for Agile Model-Based Systems Engineering - Colorado State University. (2023) Available at: https://www.engr.colostate.edu/~drherber/files/Huss2023c.pdf
9. Comparing Measured Agile Software Development Metrics Using an Agile Model-Based Software Engineering Approach versus Scrum Only - Semantic Scholar. (2023) Available at: https://pdfs.semanticscholar.org/891b/f9b38772406a16b1ed6cf135d4b5bae0ebaa.pdf
10. What is Cost of Delay? - Playbook. (2018) Available at: https://playbookteam.com/resources/blog/what-is-the-cost-of-delay
11. Ariel Dora Stern, Innovation under Regulatory Uncertainty: Evidence from Medical Technology - (2018) PMC. Available at: https://pmc.ncbi.nlm.nih.gov/articles/PMC5482535/

12. Medical Devices Compliance Guide: Costs, Budget, FDA and ISO Requirements - AJProTech. Available at: https://ajprotech.com/blog/guides-and-templates/medical-devices-compliance-guide-cost-analysis-budget-strategies-and-fda-and-iso-requirements.html
13. FDA Deficiency Cost-of-Delay Calculator - Blue Goat Cyber. Available at: https://bluegoatcyber.com/cost-of-delay
14. What is Medical Device Quality Assurance? Its Importance - Qualityze. (2026) Available at: https://www.qualityze.com/blogs/medical-device-quality-assurance
15. When Data Moves, Risk Moves with It: The Hidden Challenges of Warehousing Data. (2026) Available at: https://internalaudit360.com/when-data-moves-risk-moves-with-it-the-hidden-challenges-of-warehousing-data/
16. Integrating MES and ERP: The Missing Link in Pharma Supply Chain Digital Transformation. Available at: https://supplychainwizard.com/integrating-mes-and-erp-the-missing-link-in-pharma-supply-chain-digital-transformation/
17. Cloud Data Integration: A Guide to Real-Time Streaming Architecture - Striim. Available at: https://www.striim.com/blog/cloud-data-integration/
18. Randy Anderson and Christabel Tan, FDA Warning Letters, Consequences and Costs to the US Medical Device Industry. (2023) Available at: https://www.hilarispublisher.com/open-access/fda-warning-letters-consequences-and-costs-to-the-us-medical-device-industry-102529.html