A Governed Event Spine for Discrete Manufacturing: Traceability, Visual Management, and Enterprise Integration

Architecture, Computational Methods, and Field Evidence from the MileSoft Industry 4.0 Suite

MileSoft Engineering Research Group

MileSoft Software Technologies

https://milesoft.net/

Corresponding authors: *Deepak Daryani* (deepakdaryani@milesoft.net), *Tanuj Daryani* (tanujdaryani@milesoft.net)

Working paper — Version 1.0 — 22 August 2026

Permanent link: https://milesoft.net/research/industry-40

Abstract

Discrete manufacturers are asked to prove, after the fact, what happened to a single unit on a single shift months earlier. The evidence needed to answer that question is generated continuously by torque controllers, programmable logic controllers, barcode and RFID readers, and human buyoff stations — and is then scattered across vendor-specific log formats, spreadsheets, and paper escalation sheets. The result is a traceability gap: the data exists, but it cannot be assembled into an answer within the time an audit, a containment action, or a warranty dispute allows.

This paper presents the architecture, computational methods, and field evidence of the MileSoft Industry 4.0 software suite — a set of six interoperating modules (Torque Traceability, Andon, Material Traceability, Inventory Control, manufacturing ERP, and Annual Inventory Tagging) organised around a single governed event spine. The spine is an append-only, time-ordered event store with an identity resolver that binds every event to a unit, lot, or asset; a rule and limit layer that evaluates conformance at capture time rather than at reporting time; a KPI engine whose definitions follow ISO 22400-2; and a genealogy graph that answers forward and backward traceability queries as graph traversals rather than log joins.

We describe module-level methods spanning multi-vendor torque data acquisition and process-capability estimation, standards-conformant downtime accounting and escalation timing, lot genealogy over a directed acyclic graph, ABC-XYZ inventory segmentation, and reconciliation of a physical stock take against book stock. We anchor the work in the enterprise-control integration reference model of IEC 62264, the fastening risk classification of VDI/VDE 2862, and the capability indices of ISO 22514.

A production deployment at a Tier-1 automotive supplier, across plants at Pune and Nasik, integrated more than forty direct-current torque controllers across twelve assembly stations. The operator reports a 98% reduction in critical-fastener warranty rework, andon mean time to first response falling from six to nine minutes to under fifteen seconds, an eleven percentage-point gain in line-average overall equipment effectiveness, and IATF 16949 audit preparation compressing from three weeks to one day. We close by proposing a capability reference framework against which shop-floor traceability platforms can be evaluated.

Keywords: Industry 4.0; Manufacturing traceability; Torque traceability; Andon systems; Material genealogy; Overall equipment effectiveness; ISO 22400; IEC 62264; ISA-95; IATF 16949; VDI/VDE 2862; Process capability; OPC UA; Manufacturing ERP

1 Introduction

A discrete manufacturing line generates evidence continuously. A direct-current torque controller emits a result for every fastening it performs. A programmable logic controller emits a state change for every stop and start. A barcode or radio-frequency identification reader emits a read for every part that moves. A human buyoff station emits a signature. In principle, the record needed to answer any question about any unit already exists at the moment the unit is built.

In practice it does not survive. Each controller vendor writes its own log format. Downtime is recorded on a whiteboard or a paper escalation sheet. Genealogy is reconstructed after the fact by joining a spreadsheet of lot receipts against a spreadsheet of production output. The evidence is present but not assembled, and assembling it is a forensic exercise measured in days or weeks. This is the traceability gap: the distance between data that exists and an answer that can be produced within the window an audit, a containment action, or a warranty dispute allows.

This paper argues that the gap is architectural rather than instrumental. Adding another data-collection point to a plant that already has forty of them does not close it. What closes it is a single governed spine through which every shop-floor event passes, where identity is resolved at capture time, conformance is evaluated at capture time, and the resulting record is immutable and queryable.

1.1 The forensic bottleneck

Consider the sequence a Tier-1 supplier follows when an original equipment manufacturer returns a critical-fastener rework claim. The claim names a vehicle identification number and a date. The supplier must determine which station built that unit, which tool performed the joint, which operator was signed on, what torque and angle were achieved, and whether the tool was within calibration at the time. When each of those facts lives in a different system with a different key, the join is manual.

The cost is not only the labour of the investigation. While the join is in progress the supplier cannot bound the population at risk, so containment is necessarily conservative: more units are quarantined than were ever affected. The economic loss from over-containment routinely exceeds the loss from the original defect.

> The binding constraint is not measurement accuracy. It is the absence of a shared identity under which independently accurate measurements can be joined.

1.2 The compliance dimension

IATF 16949 requires automotive suppliers to demonstrate effective process control, and demonstration means producing evidence on demand. VDI/VDE 2862 goes further for bolted joints: it classifies joints by the consequence of failure and, for the highest categories, requires that the fastening system be capable of detecting incorrect tightenings rather than merely performing them. A plant that cannot show which joints were classified into which category, and that the tooling applied to each category met the corresponding minimum requirement, has a documentation problem that no amount of downstream inspection resolves.

Compliance therefore constrains architecture, not just process. If the classification of a joint is a property recorded in a quality document rather than a property enforced by the system that captures the fastening, the two will diverge. The design presented here treats joint classification, torque and angle limits, and escalation rules as governed master data evaluated at capture time.

1.3 Contributions

1. A governed event-spine architecture for discrete manufacturing, in which identity resolution, conformance evaluation, KPI computation, and genealogy are properties of a single shared layer rather than of six independent applications.
2. Module-level computational methods stated against published standards: multi-vendor torque acquisition and capability estimation, ISO 22400-conformant downtime and escalation accounting, genealogy as a directed acyclic graph, ABC-XYZ segmentation, and stock-take reconciliation.
3. Field evidence from a two-plant production deployment at a Tier-1 automotive supplier, with all reported metrics traced to their source.
4. An eight-dimension capability reference framework for evaluating shop-floor traceability and visual-management platforms.

2 Background and Related Work

The design draws on four established bodies of work: reference models for enterprise-control integration, the fastening-process literature and its risk classification, standardised performance measurement, and event-based traceability vocabularies.

2.1 Reference models for enterprise-control integration

IEC 62264, published in parallel as ANSI/ISA-95, organises manufacturing into levels: physical process and sensing at Levels 0 to 2, manufacturing operations management at Level 3, and business planning and logistics at Level 4. Part 1 supplies the models and terminology; Part 3 supplies an activity model for the Level 3 operations functions.

The value of the model for this work is that it names the layer at which traceability belongs. Torque results, downtime events, and material movements are Level 3 concerns: they are neither raw control signals nor business transactions. Placing the event spine explicitly at Level 3, with defined interfaces upward to enterprise systems and downward to controllers, is what allows a single architecture to serve modules that would otherwise be built as point solutions.

The strategic framing of Industry 4.0 as cyber-physical production — in which physical processes and their digital representations are continuously coupled — was set out in the acatech Industrie 4.0 working-group report and developed into a layered architecture by Lee, Bagheri and Kao. The architecture presented here is a concrete instance of that coupling for discrete assembly, with the coupling implemented as an event spine rather than as a digital twin of the machine.

2.2 Fastening: risk classification and process capability

VDI/VDE 2862 Blatt 1 classifies bolted joints in automotive applications by the consequence of failure, where failure is defined as loosening, bolt breakage, or loss of the bolt or nut. The classification determines the minimum requirement for detecting incorrect tightenings, and therefore the minimum capability of the tightening system — from a simple torque tool at the lowest category to a monitored, documented system at the highest. Blatt 2 extends the approach to plant construction, mechanical engineering, and pressure-boundary flange connections.

Where a joint is instrumented, the resulting torque and angle distributions are assessed with the capability indices defined in ISO 22514. Part 1 supplies general principles and concepts; Part 2 addresses time-dependent process models and distinguishes processes in statistical control, for which capability indices are appropriate, from those that are not, for which performance indices are the correct statistic. This distinction matters in practice: a fastening station whose mean drifts with tool temperature is not in statistical control, and reporting a capability index for it overstates confidence.

Hand torque tools used for verification are themselves governed. ISO 6789-1 specifies design and quality conformance requirements and the content of a declaration of conformance; ISO 6789-2 specifies calibration and the calculation of measurement uncertainty. A traceability record that reports a verification torque without the calibration state of the verifying tool is incomplete.

2.3 Standardised performance measurement

Overall equipment effectiveness is widely reported and inconsistently computed. ISO 22400-2 removes the ambiguity by specifying a set of key performance indicators for manufacturing operations management, each with its formula, constituent elements, time behaviour, and dimension. Availability is defined against planned production time; effectiveness against actual production time; and quality ratio against produced quantity. The purpose is comparability — two plants applying the standard to the same facts should obtain the same number.

The practical consequence for software design is that the time-state model must be captured, not inferred. A system that records only that a line stopped, without recording whether the stop fell inside or outside planned production time, cannot compute a standards-conformant availability afterwards. The event spine therefore records the time state as a first-class attribute of every stop event.

The underlying loss taxonomy originates in total productive maintenance as set out by Nakajima. ISO 22400 is best read as the normative formalisation of that tradition rather than as a replacement for it.

2.4 Event vocabularies for traceability

ISO/IEC 19987 standardises EPC Information Services, a vocabulary for visibility events that captures what happened, when, where, and why, together with the business step and disposition. Its contribution to this work is conceptual rather than protocol-level: it establishes that a traceability record is a sequence of typed events about identified objects, not a set of tables joined after the fact.

The genealogy model described in Section 4.3 follows this shape. Aggregation and transformation events — a set of components consumed into a sub-assembly, a sub-assembly consumed into a finished unit — are recorded as edges in a directed acyclic graph at the moment they occur, which is what makes forward and backward traversal a graph query rather than a reconstruction.

2.5 Incumbent tooling and its failure modes

Four failure modes recur across the deployments that motivated this work, and the architecture is a direct response to each.

- Format fragmentation. Each controller vendor exposes results through its own protocol and schema. Without a normalising acquisition layer, every additional vendor multiplies integration cost and creates a second source of truth.
- Late conformance evaluation. Where limits are applied in a reporting tool rather than at capture, the same raw result can be judged conforming by one report and non-conforming by another, depending on which version of the limit table the report used.
- Identity loss. A torque result keyed only by station and timestamp cannot be bound to a unit after the shift ends without an independent record of what was on the station at that time.
- Unversioned masters. When limits, joint classifications, and escalation rules are edited in place, the plant loses the ability to state which rule was in force when a given unit was built — which is precisely what an audit asks.

> Each failure mode follows from treating traceability as a reporting problem. Treating it as a capture-time property of a governed spine removes all four.

3 System Overview and Architecture

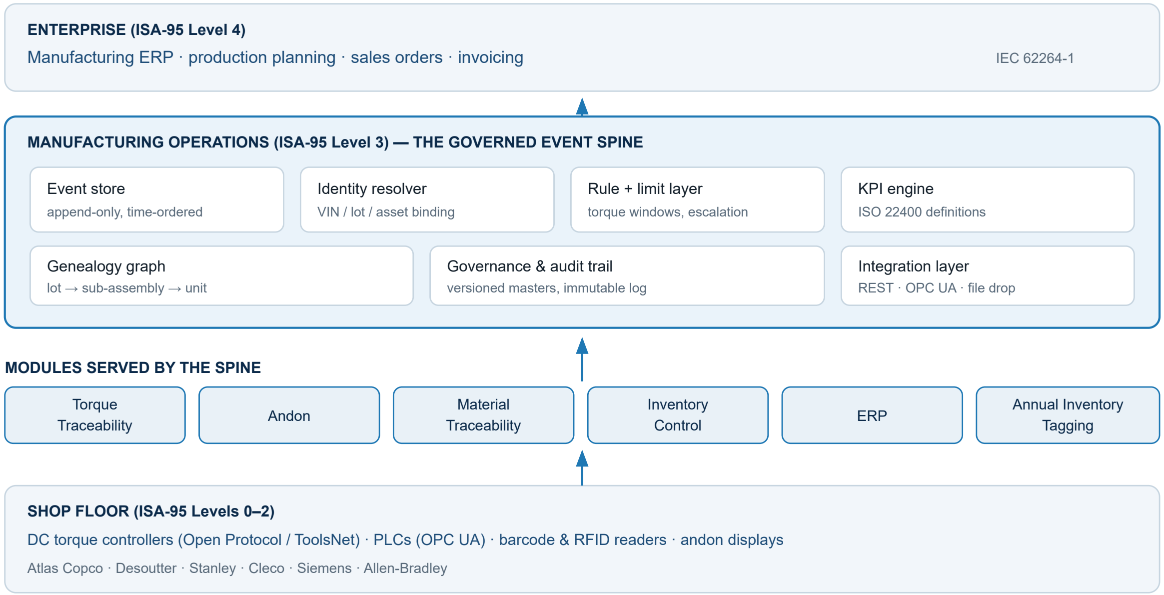

Figure 1. *Schematic.* Platform architecture. Shop-floor devices at ISA-95 Levels 0 to 2 feed a single governed event spine at Level 3 — an append-only event store, an identity resolver, a rule and limit layer, an ISO 22400 KPI engine, a genealogy graph, a governance and audit trail, and an integration layer. Six interoperating modules are served by that spine, and the spine exchanges plans and transactions with enterprise systems at Level 4.

3.1 Design goals

Six goals shaped the architecture. They are referenced by number in the module descriptions that follow.

- G1 — One identity. Every event is bound at capture time to a resolvable identity: a unit, a lot, an asset, or a location. No event is written without one.
- G2 — Conformance at capture. Limits, joint classifications, and escalation rules are evaluated when the event is written, and the verdict is stored alongside the raw value.
- G3 — Immutability. The event store is append-only. Corrections are new events that supersede, never edits that overwrite.
- G4 — Versioned governance. Every master data object — controller, limit set, joint class, escalation rule, part — is versioned, and every event references the version in force at the time.
- G5 — Standards-defined measures. Derived metrics use published definitions, principally ISO 22400-2, rather than local conventions.
- G6 — Vendor neutrality. Controller and PLC integration is isolated in adapters, so adding a vendor does not change the spine or any module.

3.2 The governed event spine

The spine has seven components. The event store is append-only and time-ordered, holding typed events with a monotonic sequence per source. The identity resolver binds each incoming event to a unit, lot, or asset using the station occupancy record — the authoritative statement of what was present at a station over a time interval — so that binding does not depend on a later join (G1).

The rule and limit layer evaluates conformance at capture: torque and angle windows, joint classification per VDI/VDE 2862, escalation thresholds, and inventory policy limits. The verdict, the rule version, and the raw value are stored together (G2, G4). The KPI engine computes ISO 22400-2 indicators from the time-state attributes carried on stop and production events (G5). The genealogy graph maintains consumption edges as a directed acyclic graph. The governance and audit trail records every master-data change as a versioned, attributed transition. The integration layer exposes REST endpoints, OPC UA connectivity, and file-drop interfaces for systems that support nothing better.

> The architectural claim is narrow and testable: because identity and conformance are resolved at capture, every question an audit asks is a query over stored facts rather than an inference over logs.

3.3 The module portfolio

Six modules are served by the spine. Torque Traceability acquires and classifies fastening results. Andon captures stops, drives escalation, and reports downtime. Material Traceability maintains genealogy from received lot to shipped unit. Inventory Control segments and governs stock. The manufacturing ERP holds production planning, sales orders, production sheets, and invoicing as the enterprise system of record. Annual Inventory Tagging reconciles a physical stock take against book stock.

The modules interoperate because they share the spine, not because they are integrated with one another. A torque non-conformance can raise an andon event, quarantine the genealogy subtree beneath the affected unit, and block the ERP dispatch of that unit, without any module holding a direct dependency on another.

3.4 Deployment, integration, and multi-plant operation

Deployment is on-premises by default, because shop-floor acquisition must continue through a wide-area network outage. Each plant runs an acquisition tier co-located with the line and a spine tier that may be shared across plants on a common campus. Controllers are integrated through vendor adapters — Open Protocol and ToolsNet for direct-current tools, OPC UA for programmable logic controllers (G6).

Where a plant loses connectivity to the shared spine, the acquisition tier buffers locally and replays on reconnection. Because the store is append-only and events carry a monotonic per-source sequence, replay is idempotent: duplicate delivery is detected and discarded rather than double-counted (G3).

3.5 Governance, versioning, and auditability

Master data is versioned rather than edited. Changing a torque window creates a new limit-set version with an effective interval; events written during the previous interval continue to reference the previous version. An auditor asking which limits applied to a unit built four months ago receives the answer from the record, not from a change log that must be interpreted.

Role-based access separates who may operate a station, who may change limits, and who may release a quarantine. Each is recorded as an attributed transition. This is the mechanism behind the reduction in audit preparation reported in Section 5: preparation collapses when the evidence is a query rather than a compilation.

4 Computational Methods by Module

This section states what each module computes and against which standard. Notation is collected in Appendix A; worked numerical examples are given in Appendix B.

4.1 Torque traceability

Acquisition normalises results from four controller families into one event schema. Each fastening event carries the achieved torque and angle, the fastening strategy, the tool and controller identity, the operator, the station, the resolved unit identity, the joint identity, and the limit-set version in force.

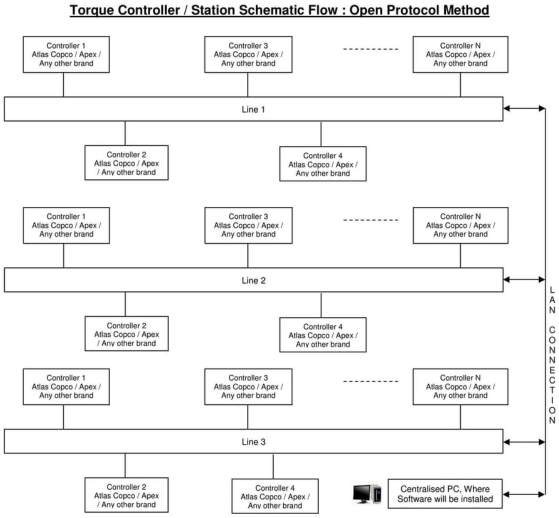

Figure 2. Torque Traceability acquisition architecture as deployed. Controllers from multiple vendors are integrated through per-vendor adapters into a single normalised result schema, so the estate can grow without changing downstream logic (design goal G6).

Each joint carries a classification derived from VDI/VDE 2862. The classification is master data with a version, not an attribute of a report, and it determines both the required detection capability and the escalation behaviour on a non-conforming result. A result outside its window is recorded as non-conforming at capture, the unit is flagged, and — for the highest joint categories — the station is inhibited until a documented rework and re-verification event is written.

Software Flow / Sequence

1. Software will Read all Controllers (50 nos) Torque data in a Daisy Chain Method
2. Compare Read Data from Torque Controller with Standard Data as per Model
3. Print data of Scanned/ Read torque data at individual Line Buy Off
4. Based upon the available/saved data Evaluate Cp & CpK (MIS) as per end user requirement
5. MIS Report (Periodical data collection download facility)
6. Read VINNo through RS Linx.

Software Features

1. Provision for Add / Delete / Modify Torque controllers

Basic commands to fetch the values from Controller: Open Protocol : MID

1. Start Command : 002000010 (Total 20 Bytes)
2. Subscribe Torque Data: 00200060000
3. End Subscribe : 00200062000
4. Continuous Read: 00209999001

Buy OFF Print format

Date	05/02/18		Time : 14:30 PM		
VIN No:	1234567890123		Body Sl No : ABCD12345		
Model / Variant : Scorpio / S4			Line	Chasis	
Result Actual OK vs Total Joints (e.g. 18 OK out of 20)					
Sl No	Joint	Set Value	Actual Value	Result	Rework
1	Steering Gear Fitment	50 - 52	55	Not OK	
2	Stab Bar Fitment	55 -70	45	Not OK	

Figure 3. Process flow from job start through fastening to buyoff. Conformance is evaluated at each fastening rather than at end of line, so the population at risk is bounded at the moment a deviation occurs (design goal G2).

Process capability is estimated per joint over a rolling window of results. Where the process is in statistical control, capability indices are reported; where it is not, ISO 22514-2 directs that performance indices be used instead, and the platform reports Pp and Ppk with the control state stated alongside.

$$Cp = (USL - LSL) / (6 \text{ sigma-hat}) \quad \text{(cp)}$$

where USL and LSL are the upper and lower torque specification limits for the joint and sigma-hat is the estimated within-subgroup standard deviation of the fastening result distribution.

$$Cpk = \min[\, (USL - mu) / (3 \text{ sigma-hat}), (mu - LSL) / (3 \text{ sigma-hat}) \,] \quad \text{(cpk)}$$

where mu is the process mean. Cpk penalises a centred-but-drifting process differently from Cp, which is why both are reported per joint rather than one alone.

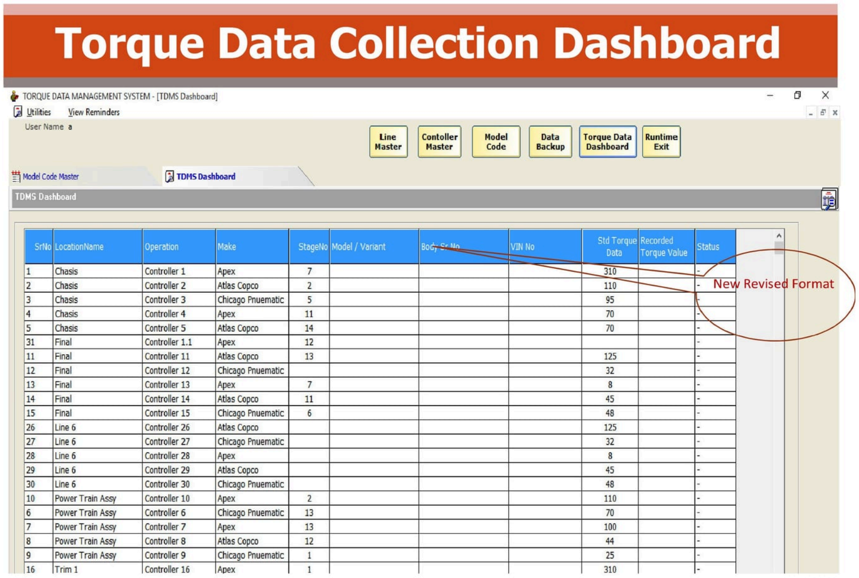

SrNo	LocationName	Operation	Make	StageNo	Model / Variant	Body Sr No	VIN No	Std Torque Data	Recorded Torque Value	Status
1	Chasis	Controller 1	Apex	7				310		-
2	Chasis	Controller 2	Atlas Copco	2				110		-
3	Chasis	Controller 3	Chicago Pnuematic	5				95		-
4	Chasis	Controller 4	Apex	11				70		-
5	Chasis	Controller 5	Atlas Copco	14				70		-
31	Final	Controller 1.1	Apex	12						-
11	Final	Controller 11	Atlas Copco	13				125		-
12	Final	Controller 12	Chicago Pnuematic					32		-
13	Final	Controller 13	Apex	7				8		-
14	Final	Controller 14	Atlas Copco	11				45		-
15	Final	Controller 15	Chicago Pnuematic	6				48		-
26	Line 6	Controller 26	Atlas Copco					125		-
27	Line 6	Controller 27	Chicago Pnuematic					32		-
28	Line 6	Controller 28	Apex					8		-
29	Line 6	Controller 29	Atlas Copco					45		-
30	Line 6	Controller 30	Chicago Pnuematic					48		-
10	Power Train Assy	Controller 10	Apex	2				110		-
6	Power Train Assy	Controller 6	Chicago Pnuematic	13				70		-
7	Power Train Assy	Controller 7	Apex	13				100		-
8	Power Train Assy	Controller 8	Atlas Copco	12				44		-
9	Power Train Assy	Controller 9	Chicago Pnuematic	1				25		-
16	Trim 1	Controller 16	Apex	1				310		-

Figure 4. Captured fastening results with their evaluated status. The raw value, the verdict, and the limit-set version are stored together, so a later change to limits cannot retroactively alter the recorded conformance of a built unit (design goals G2 and G4).

The buyoff report is generated from the same stored events rather than compiled separately. It states, per unit, every critical joint, its classification, its achieved values, its verdict, the tool and operator, and the calibration state of that tool at the time of the fastening.

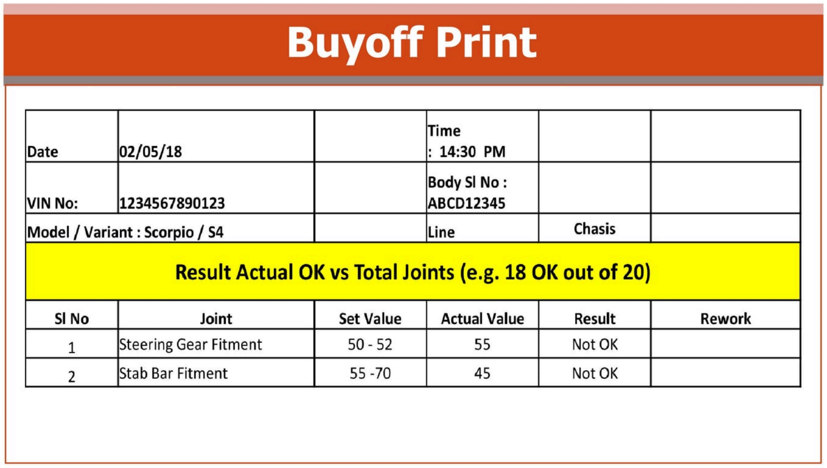

Date	02/05/18		Time : 14:30 PM		
VIN No:	1234567890123		Body Sl No : ABCD12345		
Model / Variant : Scorpio / S4			Line	Chasis	
Result Actual OK vs Total Joints (e.g. 18 OK out of 20)					
Sl No	Joint	Set Value	Actual Value	Result	Rework
1	Steering Gear Fitment	50 - 52	55	Not OK	
2	Stab Bar Fitment	55 -70	45	Not OK	

Figure 5. Unit buyoff report. Because it is generated from the event store rather than assembled from logs, the report is reproducible: regenerating it for a historical unit returns the same document.

4.2 Andon and downtime accounting

An andon event is raised by an operator, by a programmable logic controller fault, or by the rule layer in response to a non-conforming result. Each event carries its origin, the affected station, the reason code, and — critically for standards conformance — the time state in which it occurred.

Escalation is a timed state machine. The response clock starts when the event is raised, not when it is acknowledged, which is what makes mean time to first response measurable rather than self-reported. Successive tiers are paged automatically as thresholds elapse.

$$MTTR\text{-}FR = (1/n) * SUM\ over\ i\ of\ (\ t\_ack(i) - t\_raise(i)\ ) \quad \text{(mttrfr)}$$

where t_raise is the instant the event was written to the store and t_ack is the instant a responder acknowledged it. Both timestamps are system-generated; neither is entered by the responder.

Overall equipment effectiveness follows the ISO 22400-2 decomposition. Availability is operating time over planned production time; effectiveness compares the produced quantity against what the actual production time should have yielded at the reference cycle; quality ratio is good quantity over produced quantity.

$$OEE = A\ x\ E\ x\ Q \quad \text{(oee)}$$

where A = operating time / planned production time, E = (reference cycle time x produced quantity) / operating time, and Q = good quantity / produced quantity, each as defined in ISO 22400-2.

> Because the time state is captured on the stop event rather than inferred later, availability is computed from planned production time as the standard defines it. A system that infers the time state from a shift calendar will disagree with the standard whenever the calendar and reality diverge — which is exactly when the number matters.

4.3 Material traceability and genealogy

Genealogy is maintained as a directed acyclic graph whose vertices are identified objects — received lots, sub-assemblies, finished units — and whose edges are consumption relationships written at the moment of consumption. Backward traceability from a unit is a traversal toward the roots; forward traceability from a suspect lot is a traversal toward the leaves.

The economically significant quantity is not whether the traversal succeeds but how large its result is. Containment cost scales with the size of the forward closure of a suspect lot, so the design objective is to keep that closure tight by recording consumption at the unit level wherever the process permits, and at the smallest practical batch elsewhere.

$$F(lot) = \{\ units\ u : there\ exists\ a\ directed\ path\ from\ lot\ to\ u\ in\ the\ genealogy\ graph\ G\ \} \quad \text{(closure)}$$

the forward closure of a suspect lot. The population that must be contained is F(lot); every unit outside it is provably unaffected, which is what converts a conservative recall into a bounded one.

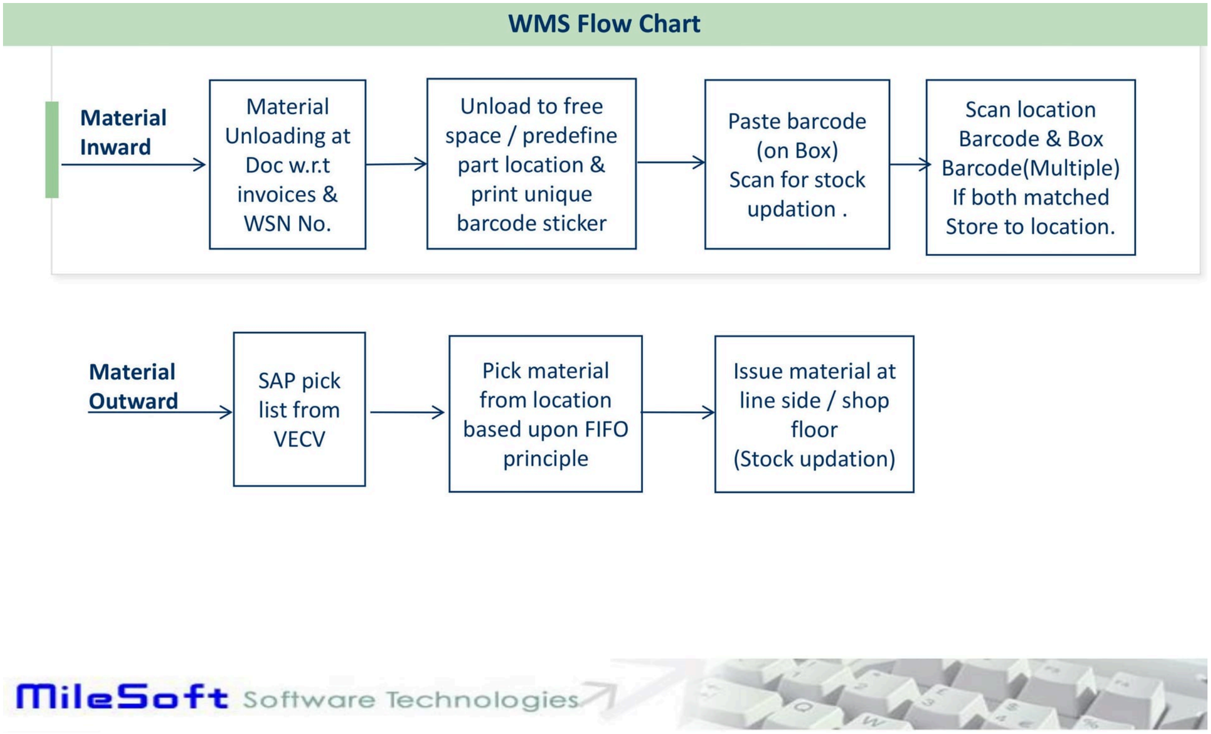

Figure 6. Material flow through receipt, storage, issue, and consumption. Each transition writes an event carrying the object identity, the location, and the business step, following the event shape standardised by ISO/IEC 19987.

4.4 Inventory control

Stock is segmented on two axes. The ABC axis ranks parts by annual consumption value and partitions the ranked list at cumulative-value thresholds. The XYZ axis classifies parts by demand variability, measured as the coefficient of variation of period demand.

$$CV(i) = sigma(i) / d\text{-}bar(i) \tag{cv}$$

where d-bar(i) is mean period demand for part i and sigma(i) its standard deviation. Low CV parts (X) admit tight reorder policies; high CV parts (Z) require either buffer stock or a make-to-order treatment.

The two axes together give nine cells, and policy is assigned per cell rather than per part. High-value, low-variability parts justify tight reorder points and frequent counting; low-value, high-variability parts justify generous buffers and infrequent counting. The reorder point follows the standard service-level formulation.

$$ROP(i) = d\text{-}bar(i) * L(i) + z(alpha) * sigma(i) * sqrt(L(i)) \tag{rop}$$

where L(i) is replenishment lead time in periods and z(alpha) is the standard normal quantile for the target cycle service level assigned to the part's ABC-XYZ cell.

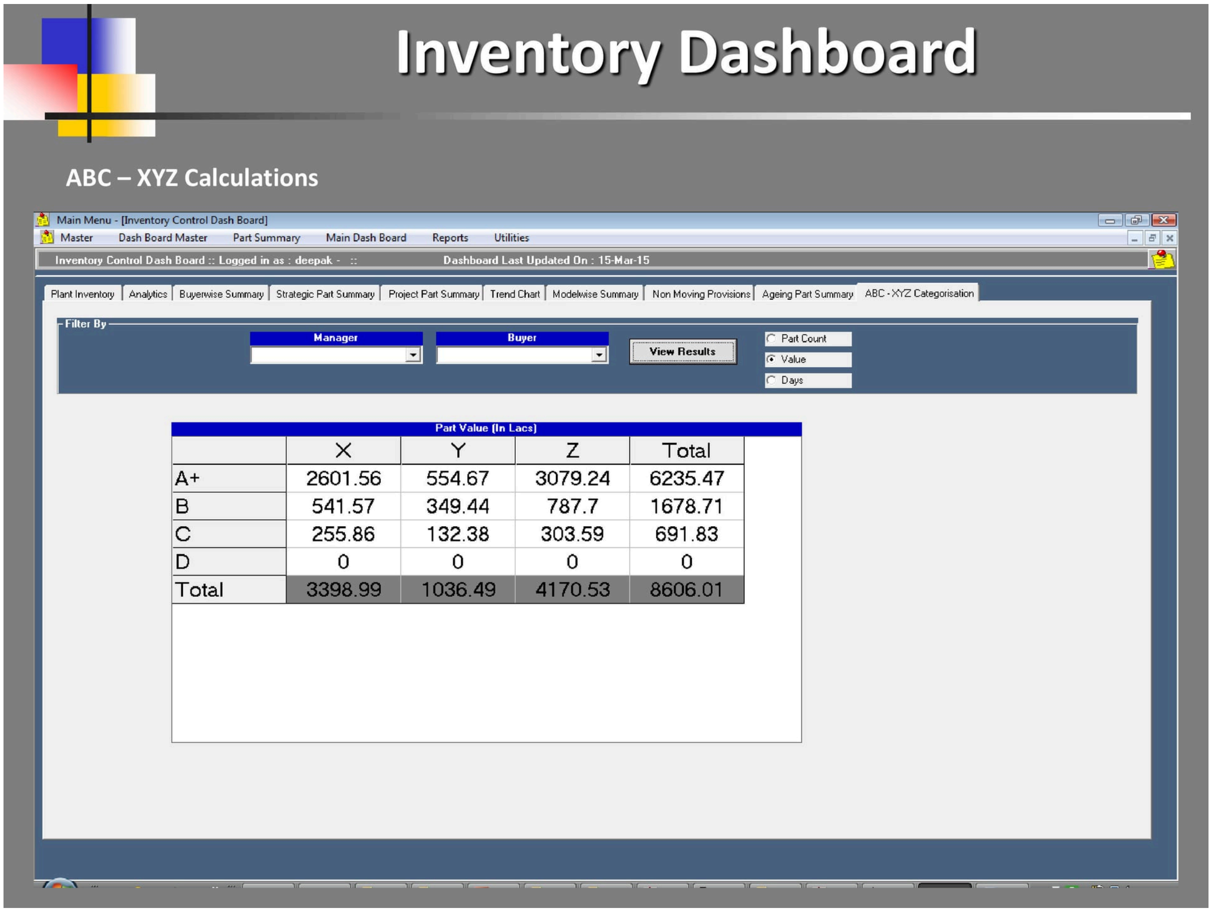

Figure 7. ABC-XYZ segmentation as deployed. Policy is assigned per cell rather than per part, which is what makes the policy set reviewable by a planner instead of being an implicit consequence of thousands of individual settings.

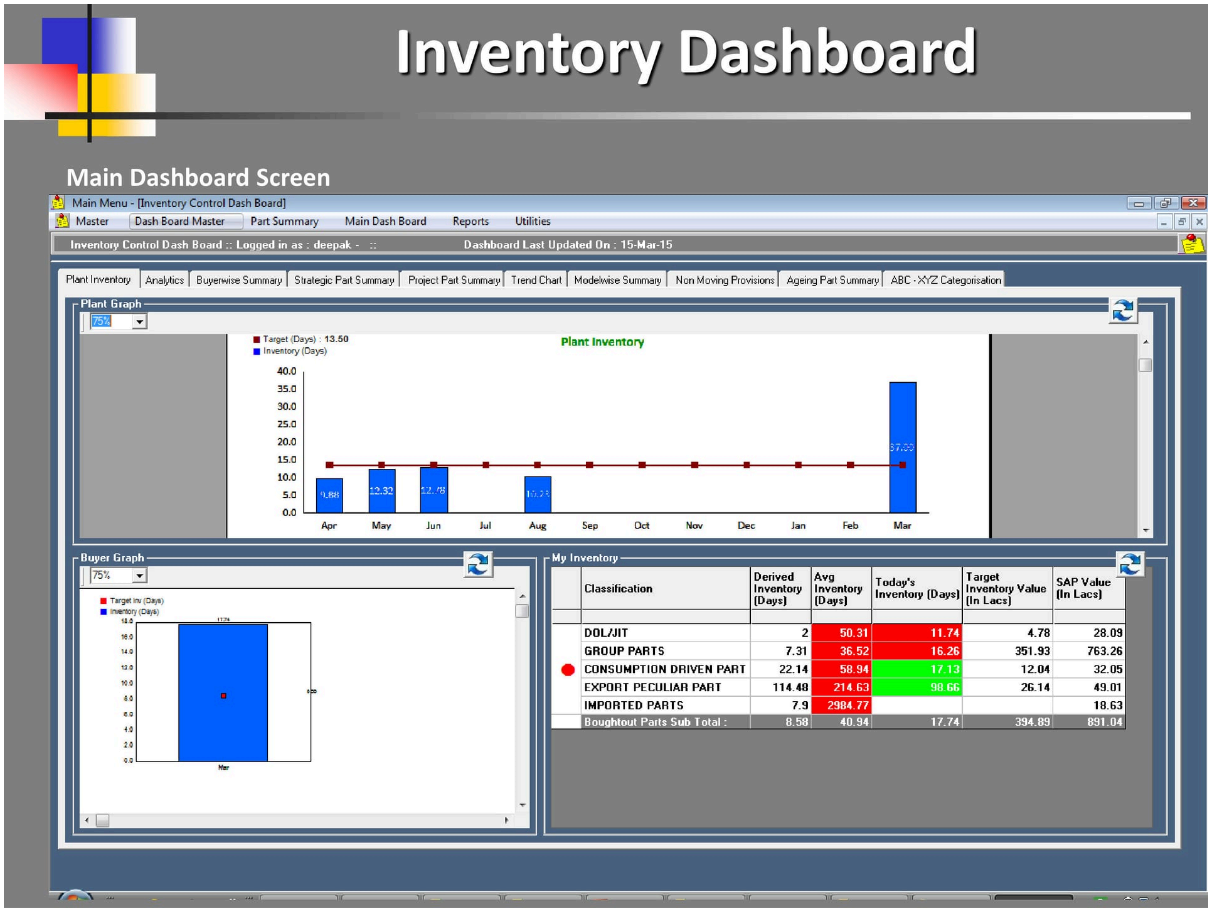

Figure 8. Inventory Control dashboard. Exceptions — non-moving stock, ageing beyond policy, parts below reorder point — are surfaced as the primary view rather than as reports that must be requested.

4.5 Manufacturing ERP

The ERP module occupies Level 4 of the IEC 62264 hierarchy: production planning, sales orders, production sheets, dispatch, and invoicing. Its relationship to the spine is deliberately narrow. Plans and orders flow down; confirmations, consumption, and quality verdicts flow up. The ERP does not read shop-floor events directly, and no module reads ERP tables directly.

This narrowness is a design choice rather than an implementation limitation. An ERP that reaches into shop-floor data becomes a second interpreter of that data, and the two interpretations diverge. Constraining the interface to plans downward and confirmations upward keeps one interpretation.

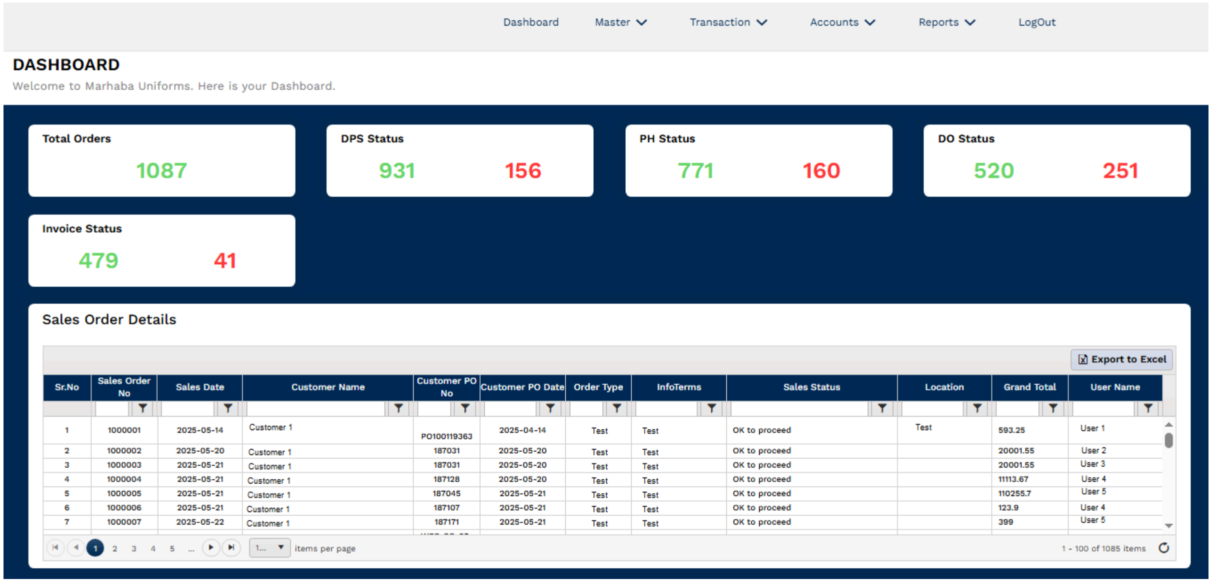

Figure 9. ERP dashboard. Production plans, sales orders, and dispatch status are held here as the enterprise system of record, while conformance and traceability evidence remains on the spine.

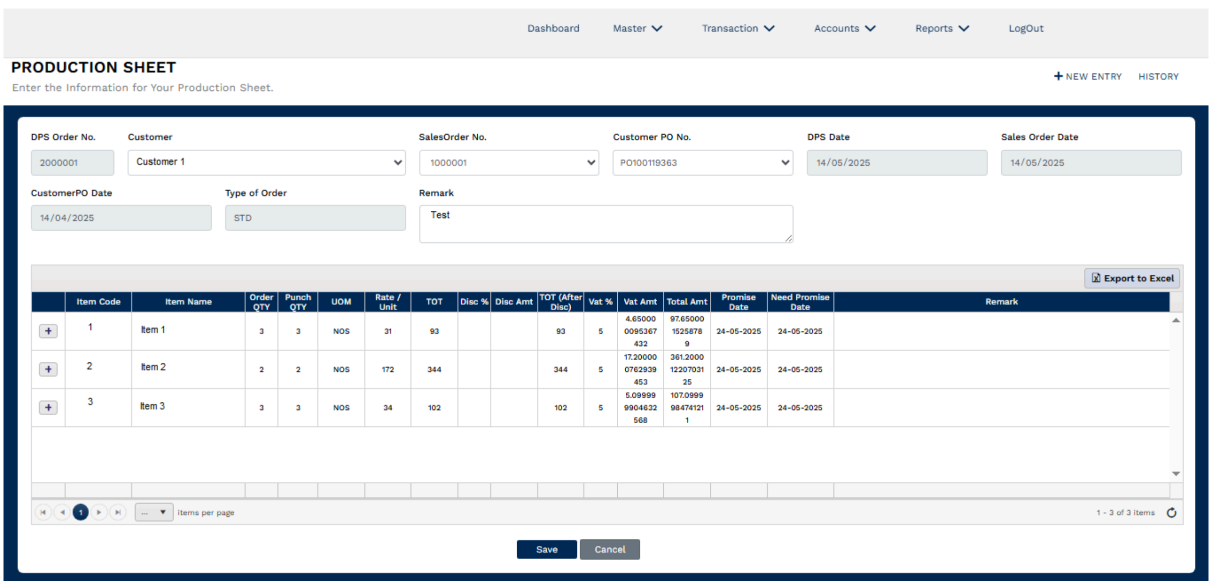

Figure 10. Production sheet. The sheet is the downward interface to the shop floor; the corresponding upward interface is the confirmation event, which carries produced and good quantities in the ISO 22400 sense.

4.6 Annual inventory tagging

A physical stock take produces counted quantities per part and location; the book position produces expected quantities. Reconciliation is the attribution of the difference. The platform records each count as an event carrying the counter, the instrument, the location, and the timestamp, so that a disputed line can be re-examined without recounting the plant.

$$IRA = (\text{number of lines } i \text{ where } |c(i) - b(i)| <= tau(i)) / N \qquad \text{(ira)}$$

inventory record accuracy on a line-count basis, where c(i) is the counted quantity, b(i) the book quantity, tau(i) the tolerance for that part, and N the number of lines counted. Line-count accuracy is reported alongside value accuracy, because the two diverge sharply when errors concentrate in low-value parts.

Variance is attributed rather than merely reported: to receipt error, issue error, scrap not booked, location error, or count error. Attribution is what turns an annual stock take from a write-off exercise into a diagnostic one, because only attributed variance identifies the process to fix.

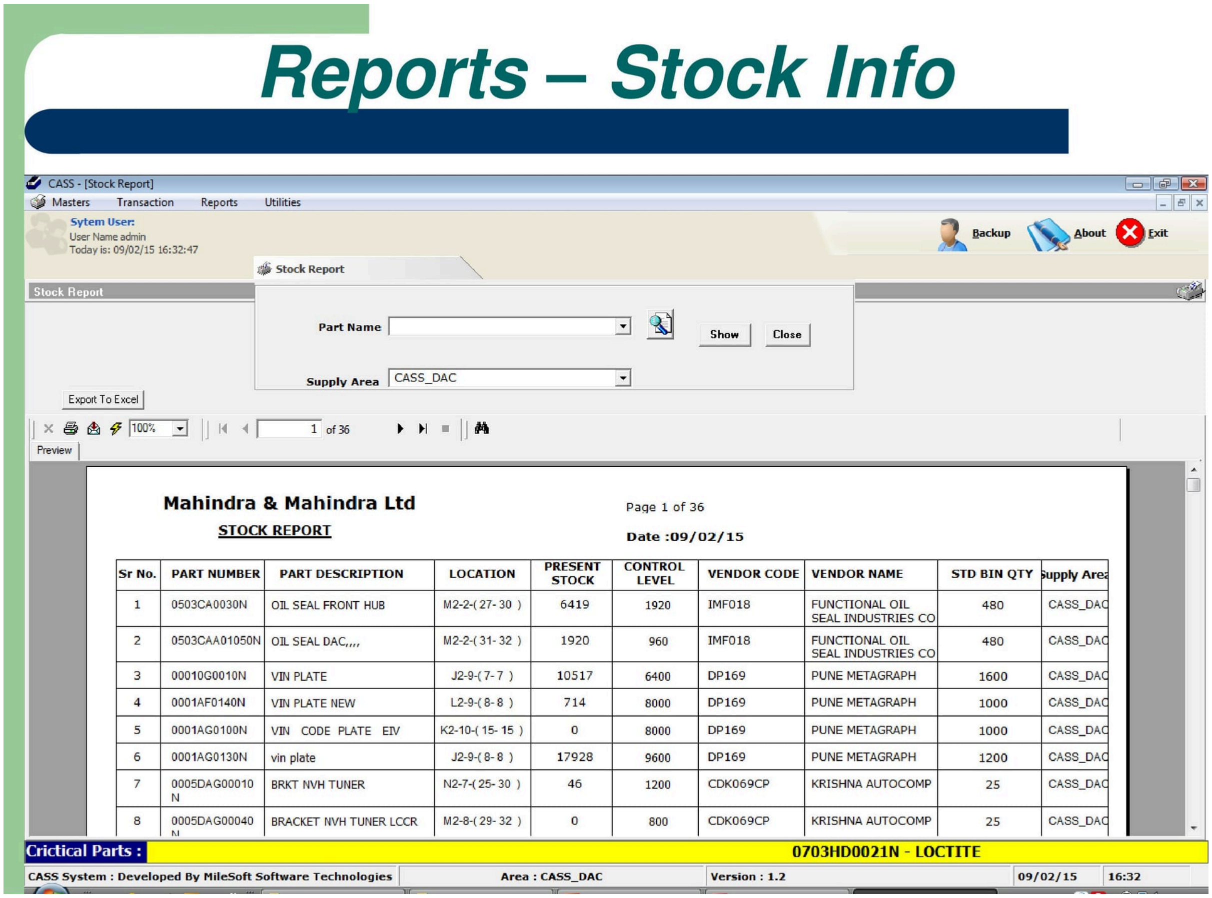

Figure 11. Annual Inventory Tagging overview. The stock take is modelled as a sequence of scoped, timestamped counting events rather than as a single plant-wide freeze, which is what allows counting to proceed alongside operations.

5 Field Evidence: A Production Deployment

The deployment described here is at a Tier-1 automotive supplier operating plants at Pune and Nasik, India. All figures in this section are operator-reported from the production system and are itemised with their provenance in the Data provenance table below.

5.1 Context and prior workflow

Before deployment, the plants operated a mixed estate of direct-current torque controllers from Atlas Copco, Desoutter, Stanley, and Cleco, each writing results in its own format. There was no single vehicle-level traceability record across the line. When a critical-fastener rework claim arrived from the original equipment manufacturer, the team needed days of log analysis to establish which station, which operator, and which time window were implicated.

Visual management was a physical stack-light with a paper escalation sheet. Mean time to first response ran between six and nine minutes. IATF 16949 audits required exhaustive evidence of process control that the existing systems could not produce in less than weeks of preparation.

5.2 What was deployed

Torque Traceability was deployed as the unified controller-to-spine bridge, integrating more than forty controllers across twelve assembly stations, with every fastening event bound to a vehicle identification number, station, operator, controller identity, and capability window. The Andon module was layered above it with programmable-logic-controller integration so that machine faults page the responsible team automatically. Material Traceability closed the genealogy loop from raw lot through sub-assembly to finished vehicle identification number. The ERP module provided the system of record behind the line.

Rollout was phased: eight weeks at the first plant, six at the second, the compression reflecting lessons carried forward rather than reduced scope.

5.3 Reported outcomes

Table 1. Operator-reported outcomes following deployment, measured against the pre-deployment baseline at the same plants. Each row is itemised with its source in the Data provenance table.

Measure	Before	After
Critical-fastener warranty rework	Baseline	98% reduction
Andon mean time to first response	6-9 minutes	Under 15 seconds
Line-average overall equipment effectiveness	Baseline	+11 percentage points
IATF 16949 audit preparation	3 weeks	1 day
Torque controllers under unified traceability	0	40+ across 12 stations

The audit-preparation figure is the clearest architectural signal in the set. Preparation time did not fall because the audit became easier; it fell because the evidence stopped needing to be compiled. Where the answer to an auditor's question is a stored fact rather than a reconstruction, preparation reduces to running the query.

The rework reduction should be read with the containment mechanism in mind. Part of the gain is defect prevention — a non-conforming fastening inhibits the station rather than passing downstream — and part is bounded containment, because the forward closure of a suspect condition is computable. The reported figure aggregates both effects and does not separate them; this is a limitation, addressed in Section 7.

6 Discussion

6.1 Capture-time conformance as the load-bearing decision

Of the six design goals, the one that produces most of the observed benefit is G2: conformance is evaluated when the event is written. Every downstream property follows from it. Containment can be bounded because non-conformance is already classified. Audit preparation collapses because the verdict is stored rather than derived. Reports become reproducible because regenerating one re-reads a verdict instead of re-applying a rule that may since have changed.

The corollary is that the value of the architecture is not proportional to the number of data points collected. A plant that captures fewer signals but resolves identity and conformance at capture will answer audit and containment questions faster than one that captures more signals and resolves neither.

6.2 Standards conformance as architecture, not reporting

ISO 22400-2 can be satisfied superficially by computing the right formula over the wrong inputs. Availability requires planned production time; if the time state is inferred from a shift calendar rather than captured on the event, the computation is standards-shaped but not standards-conformant. Placing the time state on the event is a small schema decision with a disproportionate effect on whether the resulting number is comparable across plants.

The same argument applies to VDI/VDE 2862. Joint classification recorded in a quality document is advisory; joint classification held as versioned master data and enforced by the rule layer is binding. Only the second survives contact with an auditor asking what was in force on a given date.

6.3 A capability reference framework for traceability platforms

The following eight dimensions are proposed as a vendor-independent benchmark for shop-floor traceability and visual-management platforms. They are derived from the standards landscape of Section 2 and from the failure modes catalogued in Section 2.5.

Table 2. Capability reference framework. Each dimension is stated as a question that can be answered by demonstration rather than by specification, so that an evaluation can be conducted against a running system.

Dimension	Question the platform must answer by demonstration
D1 Identity resolution	Is every captured event bound to a unit, lot, or asset at capture time, without a later join?
D2 Capture-time conformance	Is the conformance verdict stored with the raw value and the rule version that produced it?
D3 Immutability	Is the event record append-only, with corrections expressed as superseding events?
D4 Master-data versioning	Can the platform state which limits, classifications, and rules were in force on an arbitrary past date?
D5 Standards-defined measures	Are KPIs computed from captured time states using published definitions, principally ISO 22400-2?
D6 Genealogy closure	Can the forward closure of a suspect lot be computed exactly, and how tight is it?
D7 Vendor neutrality	Can a new controller or PLC vendor be added without changing the spine or any module?
D8 Degraded-mode operation	Does acquisition continue through a network outage, and is replay idempotent?

The framework is deliberately answerable by demonstration rather than by specification. Every dimension can be tested in an hour against a running system; none requires access to source code.

6.4 Generalisability

The evidence base is a single operator across two plants in one industry. The architecture is not automotive-specific — the spine makes no assumption about the nature of the unit being built — but the reported magnitudes are. Automotive assembly has an unusually high density of classified critical joints and an unusually demanding audit regime, which is precisely the environment in which capture-time conformance pays most. A plant with few critical joints and a lighter audit obligation should expect the same architectural properties with smaller measured gains.

7 Threats to Validity and Limitations

1. Single-operator evidence. All field figures come from one operator across two plants. There is no control site, and no comparison against an alternative platform deployed under matched conditions.
2. Operator-reported metrics. Figures are reported from the production system by the operator rather than independently audited. They are internally consistent and system-generated, but they are not third-party verified.

3. Confounded rollout. The modules were deployed together over a phased programme that also included process and training changes. The contribution of any single module to any single figure cannot be isolated from the others.
4. Aggregated rework reduction. As noted in Section 5.3, the 98% figure combines defect prevention and bounded containment. The two effects are not separated in the reported data.
5. No pre-registered baseline. Baselines were reconstructed from the operator's records at the start of the programme rather than measured under a protocol defined in advance.
6. Absence of cost data. The paper reports operational outcomes but not total cost of ownership, so no return-on-investment claim is made or implied.

> These limitations bound the strength of the evidence, not the description of the architecture. Sections 3 and 4 are reproducible from the text; Section 5 should be read as one operator's reported experience.

8 Future Work

- Separating prevention from containment. Instrumenting the two mechanisms independently would allow the rework reduction to be decomposed, which would in turn indicate where further investment pays.
- Multi-site comparability. Applying the ISO 22400-2 definitions across operators would test whether the standard delivers the comparability it intends once the time-state capture problem is removed.
- Closure tightness as a design metric. Reporting the expected size of the forward closure per part family would let a plant trade recording granularity against containment cost explicitly, rather than discovering the trade-off during an incident.
- Predictive capability drift. Capability indices are currently reported over a rolling window. Modelling the drift directly would allow a station to be flagged before it produces a non-conforming result rather than after.
- EPCIS interchange. The internal genealogy model follows the ISO/IEC 19987 event shape but does not yet emit conformant EPCIS documents for exchange with customers and suppliers.

9 Conclusion

The traceability gap in discrete manufacturing is not a shortage of data. It is the absence of a shared identity under which independently accurate measurements can be joined, and the absence of a conformance verdict recorded at the moment the measurement was taken. Both are architectural properties, and both can be supplied by a single governed event spine sitting at Level 3 of the IEC 62264 hierarchy.

This paper has described such a spine and the six modules it serves, stated the computational methods of each against published standards, and reported a two-plant production deployment in which critical-fastener warranty rework fell by 98%, andon first response fell below fifteen seconds, line-average overall equipment effectiveness rose eleven percentage points, and audit preparation compressed from three weeks to one day.

The capability reference framework of Section 6.3 is offered as the durable contribution. The specific platform will change; the eight questions it asks — about identity, capture-time conformance, immutability, versioning, standards-defined measures, closure, vendor neutrality, and degraded-mode operation — will remain the right questions to ask of any shop-floor traceability platform.

Appendix A Nomenclature

Table 3. Symbols and abbreviations used in this paper.

Symbol / term	Meaning
A	Availability — operating time divided by planned production time (ISO 22400-2)
E	Effectiveness — reference cycle time times produced quantity, divided by operating time
Q	Quality ratio — good quantity divided by produced quantity
OEE	Overall equipment effectiveness, the product A x E x Q
Cp, Cpk	Process capability indices for a process in statistical control (ISO 22514)
Pp, Ppk	Process performance indices, used where the process is not in statistical control
USL, LSL	Upper and lower specification limits for a joint
mu, sigma	Process mean and standard deviation of the fastening result distribution
MTTR-FR	Mean time to first response — raise to acknowledge, in the andon state machine
CV	Coefficient of variation of period demand, the XYZ segmentation axis
ROP	Reorder point for a part
L	Replenishment lead time, in demand periods
z(alpha)	Standard normal quantile for the target cycle service level
IRA	Inventory record accuracy, on a line-count basis
G	The genealogy graph — a directed acyclic graph of identified objects
F(lot)	Forward closure of a lot — the set of units reachable from it in G
VIN	Vehicle identification number
DAG	Directed acyclic graph

Appendix B Worked Numerical Examples

Appendix B.1 Fastening capability for a classified joint

A critical joint is specified at 45 Nm with a window of plus or minus 4 Nm, so LSL = 41 Nm and USL = 49 Nm. Over a rolling window of 500 fastenings the station returns a mean of 45.8 Nm with an estimated within-subgroup standard deviation of 0.95 Nm, and the process is in statistical control.

Applying Equation (cp): Cp = (49 — 41) / (6 x 0.95) = 8 / 5.70 = 1.40. Applying Equation (cpk): the upper margin is (49 — 45.8) / (3 x 0.95) = 3.2 / 2.85 = 1.12, and the lower margin is (45.8 — 41) / (3 x 0.95) = 4.8 / 2.85 = 1.68. Cpk is the smaller of the two, 1.12.

The gap between Cp = 1.40 and Cpk = 1.12 is the whole diagnostic. Spread alone would comfortably clear a 1.33 acceptance threshold; the process fails it only because the mean sits 0.8 Nm above nominal. The corrective action is a centring adjustment, not a variability-reduction programme — and reporting Cp alone would have obscured that.

Appendix B.2 Standards-conformant OEE for a shift

A shift has 480 minutes of planned production time. Recorded stops total 62 minutes inside planned production time; a further 30 minutes of planned break falls outside it and is therefore excluded from planned production time rather than counted as downtime. Operating time is 480 — 62 = 418 minutes.

Availability A = 418 / 480 = 0.871. The reference cycle time is 0.55 minutes per unit and produced quantity is 690 units, so effectiveness E = (0.55 x 690) / 418 = 379.5 / 418 = 0.908. Good quantity is 679 units, so quality ratio Q = 679 / 690 = 0.984.

Applying Equation (oee): OEE = 0.871 x 0.908 x 0.984 = 0.778, or 77.8%. Note what would happen if the 30-minute break had been treated as downtime inside planned production time, as an inferred shift calendar might: availability would fall to 388 / 480 = 0.808 and reported OEE to 72.2%. The 5.6 percentage-point difference is entirely an artefact of time-state handling, which is why ISO 22400-2 conformance requires capturing the state rather than inferring it.

Appendix B.3 Containment size from the forward closure

A supplier lot of 2,000 fasteners is found to be out of specification. Under batch-level recording the lot is consumed across four production days, and containment must cover every unit built in that window: 4 days x 690 units = 2,760 units.

Under unit-level consumption recording, the forward closure F(lot) resolves to the 1,940 units that actually received a fastener from the lot, the remaining 60 fasteners being still in stock. Containment falls from 2,760 units to 1,940, a reduction of 30% — achieved entirely by recording granularity, with no change to the physical process.

This is the trade-off named in Section 8: finer consumption recording costs capture effort at every station and repays it only during an incident. Quantifying the expected closure size per part family is what allows that trade to be made deliberately rather than discovered under pressure.

Data provenance

Every quantitative claim in this paper resolves to one of the entries below. Figures marked as modelled estimates are not deployment measurements; their assumptions are stated.

Metric	Reported value	Provenance
Critical-fastener warranty rework	98% reduction	**Operator-reported — Tier-1 Manufacturers (Pune + Nasik plants, India)** /case-studies/industry40 — Measured against the pre-deployment baseline across both plants.
Andon mean time to first response (MTTR-FR)	under 15 s (from 6-9 min)	**Operator-reported — Tier-1 Manufacturers (Pune + Nasik plants, India)** /case-studies/industry40 — PLC-integrated escalation replaced a stack-light plus paper escalation loop.
Overall equipment effectiveness, line average	+11 percentage points	**Operator-reported — Tier-1 Manufacturers (Pune + Nasik plants, India)** /case-studies/industry40
IATF 16949 audit preparation time	1 day (from 3 weeks)	**Operator-reported — Tier-1 Manufacturers (Pune + Nasik plants, India)** /case-studies/industry40
DC torque controllers integrated	40+ across 12 assembly stations	**Operator-reported — Tier-1 Manufacturers (Pune + Nasik plants, India)** /case-studies/industry40 — Mixed estate: Atlas Copco, Desoutter, Stanley, Cleco.
Phased rollout duration	8 weeks (plant 1), 6 weeks (plant 2)	**Operator-reported — Tier-1 Manufacturers (Pune + Nasik plants, India)** /case-studies/industry40
Mean time to repair reduction attributed to Andon	60-80%	**Product specification** /products/andon

References

[1] Kagermann, H., Wahlster, W., and Helbig, J. (2013). Securing the future of German manufacturing industry: Recommendations for implementing the strategic initiative INDUSTRIE 4.0. Final report of the Industrie 4.0 Working Group, acatech - National Academy of Science and Engineering. https://www.acatech.de/wp-content/uploads/2018/03/Final_report__Industrie_4.0_accessible.pdf

[2] Lee, J., Bagheri, B., and Kao, H.-A. (2015). A cyber-physical systems architecture for Industry 4.0-based manufacturing systems. Manufacturing Letters, 3, 18-23.

[3] Nakajima, S. (1988). Introduction to TPM: Total Productive Maintenance. Productivity Press, Cambridge, MA.

[4] International Electrotechnical Commission / International Society of Automation (2025). Enterprise-control system integration - Part 1: Models and terminology. IEC 62264-1; ANSI/ISA-95.00.01-2025 (IEC 62264-1 Mod).

[5] International Society of Automation (2013). Enterprise-control system integration - Part 3: Activity models of manufacturing operations management. ANSI/ISA-95.00.03-2013 (IEC 62264-3 Modified).

[6] International Organization for Standardization (2014). Automation systems and integration - Key performance indicators (KPIs) for manufacturing operations management - Part 1: Overview, concepts and terminology. ISO 22400-1:2014. https://www.iso.org/standard/56847.html

[7] International Organization for Standardization (2014). Automation systems and integration - Key performance indicators (KPIs) for manufacturing operations management - Part 2: Definitions and descriptions. ISO 22400-2:2014. https://www.iso.org/standard/54497.html

[8] International Automotive Task Force (2016). Quality management system requirements for automotive production and relevant service parts organizations. IATF 16949:2016, first edition, 1 October 2016 (superseding ISO/TS 16949).

[9] Verein Deutscher Ingenieure / Verband der Elektrotechnik (2025). Minimum restrictions for application of fastening systems and tools - Applications in the automotive industry. VDI/VDE 2862 Blatt 1.

[10] Verein Deutscher Ingenieure / Verband der Elektrotechnik (2019). Minimum requirements for application of fastening systems and tools - Applications in plant construction, mechanical engineering, equipment manufacturing and for flange connections in components under pressure boundary. VDI/VDE 2862 Blatt 2.

[11] International Organization for Standardization (2014). Statistical methods in process management - Capability and performance - Part 1: General principles and concepts. ISO 22514-1:2014. https://www.iso.org/standard/64135.html

[12] International Organization for Standardization (2017). Statistical methods in process management - Capability and performance - Part 2: Process capability and performance of time-dependent process models. ISO 22514-2:2017. https://www.iso.org/standard/71617.html

[13] International Organization for Standardization (2017). Assembly tools for screws and nuts - Hand torque tools - Part 1: Requirements and methods for design conformance testing and quality conformance testing: minimum requirements for declaration of conformance. ISO 6789-1:2017. https://www.iso.org/standard/62549.html

[14] International Organization for Standardization (2017). Assembly tools for screws and nuts - Hand torque tools - Part 2: Requirements for calibration and determination of measurement uncertainty. ISO 6789-2:2017. https://www.iso.org/standard/62550.html

[15] International Electrotechnical Commission (2025). OPC unified architecture - Part 1: Overview and concepts. IEC 62541-1:2025. https://webstore.iec.ch/en/publication/81513

[16] International Electrotechnical Commission (2020). OPC unified architecture - Part 2: Security model. IEC 62541-2:2020.

[17] International Electrotechnical Commission (2020). OPC unified architecture - Part 4: Services. IEC 62541-4:2020.

[18] International Electrotechnical Commission (2020). OPC unified architecture - Part 14: PubSub. IEC 62541-14.

[19] International Organization for Standardization / International Electrotechnical Commission (2024). Information technology - EPC Information Services (EPCIS) Standard, version 2.0. ISO/IEC 19987:2024. https://www.iso.org/standard/85557.html

[20] International Organization for Standardization / International Electrotechnical Commission (2015). Information technology - EPC Information Services (EPCIS) Standard, version 1.1. ISO/IEC 19987:2015. https://www.iso.org/standard/66796.html