A Digital Pedigree for Every Joint: Multi-Vendor Torque Capture, VIN Binding, and Capability at the Fastening

Architecture, Computational Methods, and Field Evidence from MileSoft Torque Traceability

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 — 23 August 2026

Permanent link: https://milesoft.net/research/products/torque-traceability

Abstract

The joints that appear in automotive recalls are a short and predictable list: brakes, steering, suspension, airbag mounts. They are also the joints a supplier is least able to defend, because the evidence that a particular fastening on a particular vehicle met specification is distributed across controllers from four or five vendors, each writing results in its own format, none of them bound to the vehicle that was being built.

This paper presents the architecture and computational methods of MileSoft Torque Traceability, which integrates every major controller brand into a single daisy chain over Open Protocol, ToolsNet or TCP/IP, evaluates each fastening against its stored joint specification at the moment it occurs, and binds the result to a vehicle identification number obtained through barcode, RFID or a programmable-logic-controller handshake.

We state the capability estimation performed per joint, per shift and per line — Cp, Cpk, Pp and Ppk — and the control-state distinction that governs which pair may legitimately be reported; the identity-binding condition that makes a fastening record admissible; and the traceability query structure that produces a per-vehicle history. The work is anchored in IATF 16949 for evidence-of-process-control obligations, VDI/VDE 2862 for joint risk classification, ISO 22514 for capability and performance indices, and ISO 6789 for the calibration state of verification tooling.

The product's published record reports that audit preparation falls from weeks to one day once the traceability spine is live, that containment requests are answered in minutes rather than days, and that traceability queries resolve in under a minute. A four-module deployment at a Tier-1 automotive supplier, in which this module was one of four, reported a 98% reduction in critical-fastener warranty rework; that figure belongs to the programme and is not attributed to this module alone.

Keywords: Torque traceability software; IATF 16949 audit trail; VDI/VDE 2862 joint classification; Cp Cpk process capability; VIN linking; Open Protocol ToolsNet; Multi-brand controller integration; Buyoff report; Statistical process control; Critical fastener; Automotive assembly quality; Torque drift detection

1 Introduction

A critical-fastener warranty claim arrives from an original equipment manufacturer naming a vehicle identification number and a date. The supplier must establish 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.

Every one of those facts was measured. None of them, in a conventional installation, was written down together. The torque controller recorded a result keyed to itself; the line recorded a build sequence; the calibration log lives in a quality system. Assembling them into an answer is a forensic exercise, and the claim has a response deadline that the exercise does not respect.

This paper describes a system built so that the answer is a stored record rather than a reconstruction. Its two load-bearing decisions are that every controller writes into one normalised schema regardless of vendor, and that every fastening is bound to a vehicle identity at the moment it occurs.

1.1 The cost of not knowing

The direct cost of a forensic investigation is the labour it consumes, and that is the smaller half. The larger half is that while the investigation runs, the supplier cannot bound the population at risk. Containment must therefore be conservative, and a conservative containment quarantines vehicles that were never affected.

This is why traceability query time is an economic variable rather than a convenience. The product's own record states that containment requests are answered in minutes rather than days and that traceability queries resolve in under a minute. The value of that is not the time saved answering; it is the units not quarantined.

> A supplier who cannot bound the affected population pays for the units they cannot exclude. Query speed is containment size by another name.

1.2 Contributions

1. A normalising acquisition layer that integrates five controller vendors over Open Protocol, ToolsNet and TCP/IP into one fastening-event schema.
2. An identity-binding condition making a fastening record admissible only when it resolves to a vehicle identification number at capture time.
3. Capability estimation per joint, per shift and per line, with the control-state distinction that determines whether capability or performance indices may be reported.

4. An eight-dimension capability reference framework for torque traceability systems.

2 Background and Related Work

Three regimes bound what a torque traceability system must be able to state: the quality-management obligation, the fastening risk classification, and the statistical treatment of the results it captures.

2.1 The evidence obligation

IATF 16949 requires automotive suppliers to demonstrate effective process control, and demonstration means producing evidence on request. It further requires that error-proofing devices be verified and that a documented reaction plan exist for the case where a device fails.

Both requirements have architectural consequences rather than procedural ones. Verification must be scheduled and recorded per device, which means the system must know its own instrument estate. And a reaction plan must be enforceable rather than advisory — the station has to know what to do when its own sensor becomes suspect, which is a runtime behaviour rather than a document.

2.2 Which joints must be detected, not merely tightened

VDI/VDE 2862 Blatt 1 classifies bolted joints in automotive applications by the consequence of failure — 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, ranging from a simple torque tool at the lowest category to a monitored and documented system at the highest.

The consequence for software is that joint classification must be governed master data enforced by the capture layer, not an attribute recorded in a quality document. Where classification lives in a document, it and the tooling actually applied will diverge, and the divergence is discovered during an audit rather than before one.

2.3 Capability, performance, and the control state

Fastening results form a distribution, and ISO 22514 provides the indices that describe it. Part 1 supplies general principles and concepts; Part 2 addresses time-dependent process models and draws the distinction that matters operationally: processes in statistical control, for which capability indices Cp and Cpk are appropriate, and processes that are not, for which performance indices Pp and Ppk are the correct statistic.

This is not a technicality. A fastening station whose mean drifts with tool temperature across a shift is not in statistical control, and reporting Cpk for it overstates confidence in a way that is invisible in the number itself. A system that computes both pairs and states the control state alongside them is reporting something an auditor can act on.

Verification tooling carries its own governance. ISO 6789-1 specifies design and quality conformance requirements and the content of a declaration of conformance for hand torque tools; ISO 6789-2 specifies calibration and the determination of measurement uncertainty. A traceability record quoting a verification torque without the calibration state of the verifying tool is incomplete.

> Two numbers with the same name mean different things depending on a property of the process that neither number shows. Reporting the control state is what makes capability reporting honest.

2.4 Failure modes of conventional torque data collection

- Vendor-fragmented logs. Each controller writes its own format, so the audit trail has as many shapes as the plant has tool brands.
- Station-keyed results. A result identified by station and timestamp cannot be bound to a vehicle after the shift without an independent occupancy record.
- Late limit application. Where windows are applied in a reporting tool rather than at capture, the same raw result can be judged conforming or not depending on which limit version the report used.
- Capability without control state. Cpk is reported for processes that are not in statistical control, overstating confidence.
- Unverified verification. Hand-torque audit results are recorded without the calibration state of the tool that produced them.

3 System Overview

The system is organised as an acquisition layer over a mixed controller estate, a governed master-data layer describing lines, controllers, model codes and joint specifications, a capture-time evaluation layer, and a reporting layer that produces buyoff certificates and traceability queries from the same stored events.

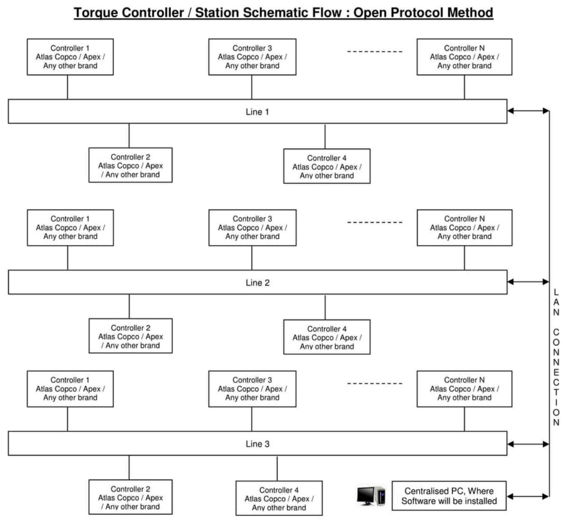

Figure 1. The acquisition architecture as deployed. Controllers from several vendors are connected in the same daisy chain and normalised into one fastening-event schema. This is what allows a plant to add a tool brand without adding a second audit trail.

3.1 The acquisition layer

Every major controller brand is integrated — Atlas Copco, Desoutter, Stanley, Cleco and Ingersoll Rand — over Open Protocol, ToolsNet or direct TCP/IP, with vendor adapters isolating the protocol differences. The design objective is that adding a vendor changes an adapter and nothing else.

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 2. The process sequence. Controller data is read, compared against the joint's stored specification, and a pass or fail returned immediately — with rework routing where the result falls outside its window. Evaluation happens here rather than in a later report.

Fault signalling runs in the reverse direction. Where a fastener is rejected, the system signals the programmable logic controller, so the line's own interlocks participate in the reaction rather than depending on an operator noticing a screen.

3.2 Governed master data

Four master records govern what the capture layer will accept: the line, the controller, the model code, and the joint specification with its torque and angle window and its VDI/VDE 2862 classification. These are the objects an auditor's question resolves against.

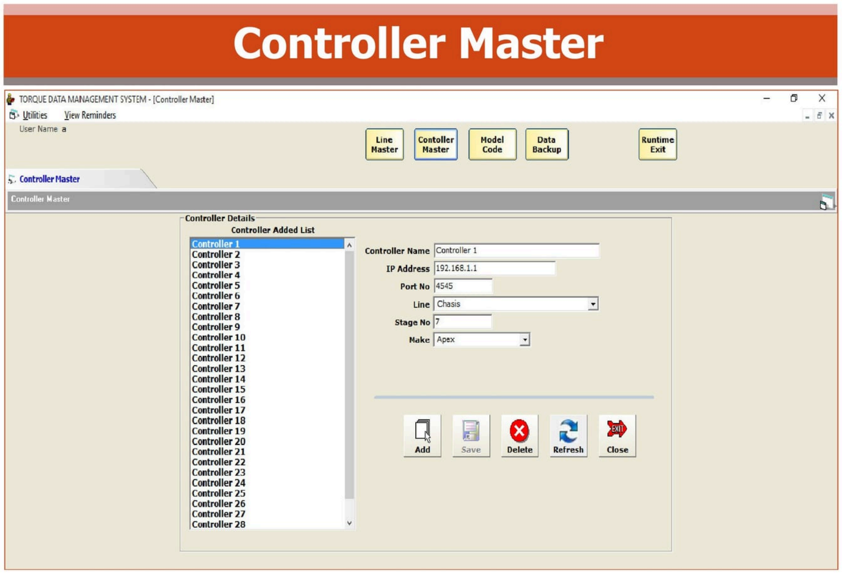

Figure 3. Controller master. Holding the instrument estate as governed data is what makes the IATF 16949 verification obligation enforceable — the system knows which devices exist, and therefore which are overdue.

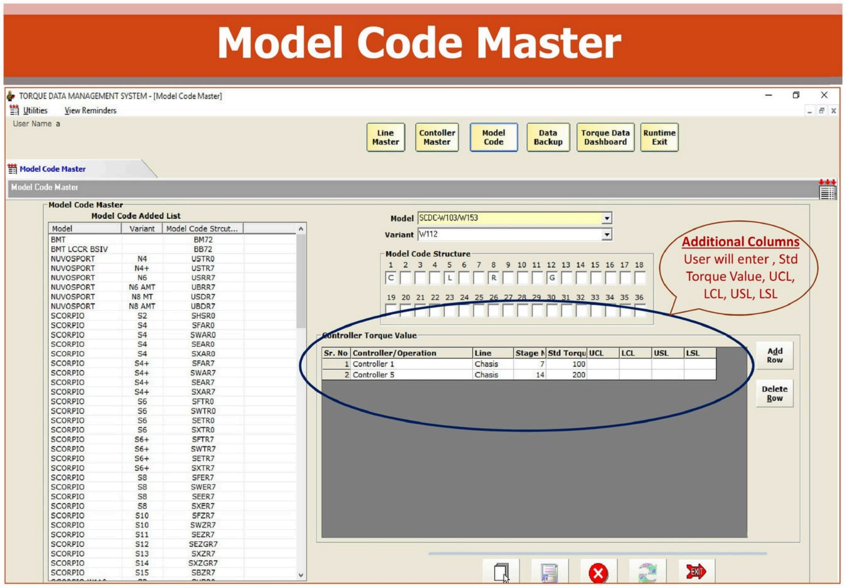

Figure 4. Model code master. A vehicle model determines which joints exist and which specification each carries, so model code is the bridge between the vehicle identity captured at the station and the windows the fastening is evaluated against.

3.3 Capture and identity binding

Vehicle identity enters through barcode scanners, RFID or a programmable-logic-controller handshake — including Allen-Bradley over RS Linx — and is resolved before the fastening is written. A result that cannot be bound to a vehicle is not a weaker record; it is not a record at all for traceability purposes.

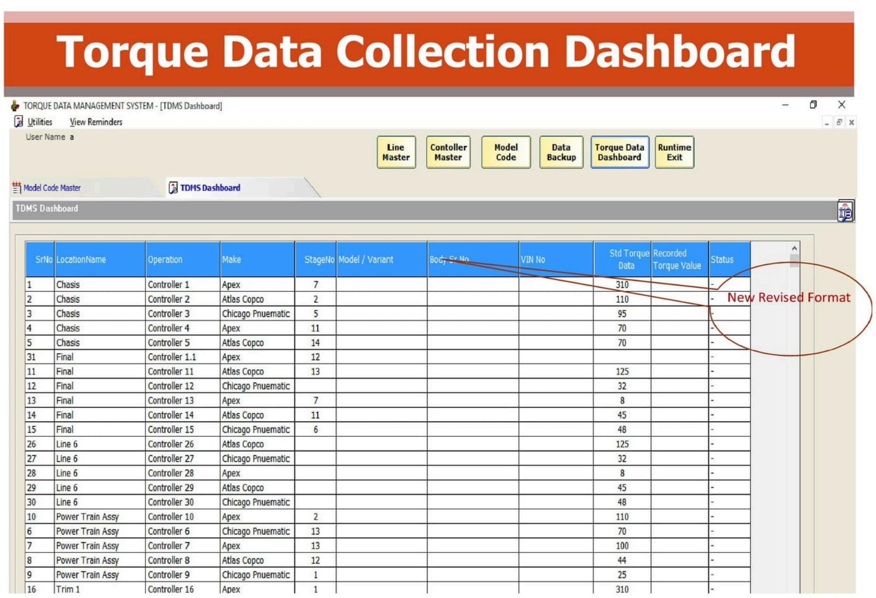

Figure 5. Live capture. Torque values, controller status and the evaluated compliance result appear together as fastenings occur. Storing the raw value, the verdict and the specification version in one record is what makes a historical buyoff reproducible.

3.4 Buyoff and the searchable archive

Station buyoff certificates print automatically at shift end or per vehicle, listing every critical joint with its achieved values, its verdict, the tool and the operator. Because they are generated from stored events, regenerating one for a historical vehicle reproduces the document rather than rebuilding it.

Buyoff Print

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 6. A unit buyoff certificate. This document is the artefact a warranty dispute turns on, and its value depends entirely on being reproducible — which requires that the verdict and the specification version were stored with the raw value, not recomputed at print time.

The archive is searchable by vehicle identification number, which is the access path a warranty investigation, a recall defence or an IATF 16949 audit actually uses. Section 4.3 states the query structure this supports.

4 Computational Methods

Notation is collected in Appendix A; worked numerical examples in Appendix B.

4.1 The admissibility condition

A fastening event is admissible as traceability evidence only when four things resolve at capture time: the vehicle identity, the joint identity, the specification version in force, and the calibration state of the tool.

$$\textit{Admissible(event) = 1 only if: VIN resolved AND joint is in the governed joint set AND a specification version is bound AND the tool's time since calibration is within its interval} \tag{admissible}$$

each factor is an indicator that is 1 when the condition holds and 0 otherwise, so the product is 1 only when all four hold. An event failing any factor is written to the store but flagged, rather than being silently accepted or silently dropped.

Writing rather than dropping an inadmissible event is deliberate. A discarded event leaves a gap that looks identical to a fastening that never happened, and the difference between those two matters enormously in a dispute.

4.2 Capability, performance, and which to report

Capability indices describe the relationship between the specification window and the observed distribution. Cp treats spread alone; Cpk penalises a mean that is off-centre.

$$Cp = (USL - LSL) / (6 \times \text{sigma-within}),\ Cpk = \min[\ (USL - mu) / (3 \times \text{sigma-within}),\ (mu - LSL) / (3 \times \text{sigma-within})\ ] \quad \text{(cpk)}$$

sigma-within is the within-subgroup standard deviation, which estimates short-term variation only. USL and LSL are the joint's upper and lower torque specification limits.

Performance indices use the overall standard deviation instead, which includes drift between subgroups. The two coincide when the process is in statistical control and diverge when it is not — and the size of the divergence is itself diagnostic.

$$Pp = (USL - LSL) / (6 \times \text{sigma-overall}), \text{ and the ratio } Cpk / Ppk \text{ is approximately sigma-overall} / \text{sigma-within} \quad \text{(ppk)}$$

a ratio near 1 indicates a stable process; a ratio well above 1 indicates drift between subgroups that the within-subgroup estimate does not see. ISO 22514-2 directs that performance indices be reported where the process is not in control.

> Reporting Cpk alone for a drifting station is the most common way a capability number flatters a process. The ratio to Ppk is the cheapest available check, and it costs nothing once both are computed.

4.3 Traceability query structure

Two query directions matter, and they have different shapes. Backward traceability from a vehicle returns every fastening event bound to it. Forward traceability from a suspect condition — a tool that drifted, a specification version that was wrong, a shift — returns every vehicle affected.

$$\text{Backward(VIN)} = \{ \text{ events e bound to that VIN } \},\ \text{Forward(condition)} = \{ \text{ distinct VINs of events satisfying the condition } \} \quad \text{(query)}$$

the forward set is the containment population. Its size is the economic quantity: every vehicle outside it is provably unaffected, and every vehicle inside it must be handled.

Both are index lookups rather than joins because identity was bound at capture. This is the mechanism behind the product's reported sub-minute query time — not query optimisation, but the absence of a reconstruction step.

4.4 Drift detection before non-conformance

Statistical-process-control dashboards flag torque drift before it produces an out-of-window result, which is the difference between a maintenance action and a containment action. The trigger is a shift in the subgroup mean relative to the specification window.

$$Z(k) = [\text{ subgroup mean at } k \text{ - established process mean }] / [\text{ sigma-within } / \text{ sqrt(subgroup size) }] \quad \text{(drift)}$$

successive Z values crossing a configured threshold indicate the mean has moved even while every individual result remains inside the window. The station is still producing conforming joints, and it is on a trajectory that will stop doing so.

5 Reported Outcomes and Field Evidence

Three kinds of number are separated, as in the other papers in this series.

5.1 Figures published for this module

Table 1. Outcomes published in this product's own record. These are the module's stated results, not a controlled study.

Measure	Reported	Where stated
IATF 16949 audit preparation	Weeks to one day once the spine is live	Torque Traceability blog post
Response to an OEM containment request	Minutes, not days	Torque Traceability blog post
Per-VIN torque history retrieval	Within an hour	Torque Traceability blog post
Forward and backward traceability query	Sub-minute	Torque Traceability blog post
Critical-joint compliance	100% of joints meet specification	Product page benefits

The audit-preparation figure is the one with an architectural explanation rather than an operational one. Preparation did not become faster; it became unnecessary, because the audit trail generates itself continuously for every vehicle that leaves the line. What used to be a compilation is now a query.

The 100% compliance claim needs reading precisely. It states that every joint is evaluated against its specification and non-conforming results are caught, not that no non-conforming fastening ever occurs. The system's function is to ensure none passes undetected, which is a different and more defensible proposition.

5.2 A deployment in which this module was one of four

MileSoft's published Industry 4.0 case study describes a deployment at a Tier-1 automotive supplier across plants at Pune and Nasik, in which Torque Traceability was installed alongside the Andon System, Material Traceability and the manufacturing ERP. More than forty controllers across twelve assembly stations were brought under unified traceability. The programme reported a 98% reduction in critical-fastener warranty rework and IATF 16949 audit preparation falling from three weeks to one day.

> Those figures belong to the four-module programme. This module's individual contribution was not isolated, and this paper does not claim them for it.

The audit figure is nonetheless the one most plausibly attributable here, because the case study and this module's own published record report the same outcome independently — three weeks to one day in the deployment, weeks to one day in the product record. Agreement between two sources is not isolation, but it is more than a single assertion.

5.3 Modelled error avoidance

Table 2. Published model parameters and the default-case result. Every input is adjustable; the result is a modelled estimate rather than a measurement.

Parameter	Default
Torqued joints per month	200,000
Baseline torque error and rework rate	0.5%
Reduction attributed to the module	60%
Fastening errors avoided per year	7,200 of 12,000

The 60% reduction is the model's central assumption and the least defensible part of it. Unlike a picking interlock, which makes a wrong action physically impossible, torque traceability detects rather than prevents — so its effect depends on what the plant does with a detection. The figure is best read as describing a plant that acts on its alerts.

6 Discussion

6.1 Identity binding is the decision everything else rests on

Of the design choices described here, binding vehicle identity at capture produces the rest. Sub-minute queries follow from it, because a bound record is an index lookup rather than a reconstruction. Bounded containment follows from it, because the forward set is computable. Reproducible buyoff follows from it, because a historical document can be regenerated against the identity it was written for.

The corollary is that improving any downstream capability without binding identity does not work. A faster reporting engine over station-keyed results still cannot answer the question an auditor asks, because the join it needs does not exist in the data.

6.2 Detection is weaker than prevention, and should be described as such

Torque traceability detects a non-conforming fastening; it does not make one impossible. That is a real limitation relative to a physical interlock, and it is worth stating plainly in a paper that also argues for the system's value.

What the architecture provides is that detection is immediate rather than downstream, that the unit is flagged rather than passing, and that for the highest joint categories the station can be inhibited pending a documented rework. The gap between detection and prevention narrows to the extent that the reaction is enforced by the system rather than left to a person — which is why the reaction plan being a runtime state rather than a document is the load-bearing part.

6.3 A capability reference framework for torque traceability systems

Table 3. Capability reference framework. Each dimension is answerable by demonstration against a running line, without access to source code.

Dimension	Question the system must answer by demonstration
D1 Vendor neutrality	Add a controller from a new brand. Does anything beyond an adapter change?
D2 Identity at capture	Is a fastening admissible without a resolved VIN, or is it flagged?
D3 Capture-time evaluation	Is the verdict stored with the raw value and the specification version that produced it?
D4 Version recovery	For a vehicle built a year ago, can the system state which torque window applied?
D5 Control state	Are Cp/Cpk and Pp/Ppk both computed, and is the control state reported alongside them?
D6 Forward closure	Name a suspect tool and a date range. How long to return the exact affected VIN list?
D7 Device verification	What happens when a controller is overdue for verification? Demonstrate the state.
D8 Reproducible buyoff	Regenerate a buyoff certificate from six months ago. Is it identical to the original?

D8 is the strongest single test. A buyoff that regenerates identically proves the verdict and the specification version were stored rather than recomputed — which is what makes the document evidence rather than an assertion.

6.4 Generalisability

The architecture generalises to any assembly process producing per-unit measured results against a governed specification — fastening, dispensing, welding, press-fit. It is automotive-specific only in its regulatory framing; the same structure serves AS9100 aerospace requirements with a different classification scheme. The reported figures come from automotive assembly with a high density of classified critical joints, and a plant with few such joints should expect the same architectural properties with smaller measured effects.

7 Threats to Validity and Limitations

1. Vendor-reported figures. Every claim in Section 5.1 is published by the vendor from its own record, without sample sizes, distributions or independent audit.

2. Confounded deployment result. The 98% rework reduction belongs to a four-module programme with concurrent process change; this module's contribution was not isolated.
3. Detection, not prevention. The system catches non-conforming fastenings; it does not make them impossible. Its effect therefore depends on the plant's reaction, which the reported figures do not separate from the software.
4. The 60% modelled reduction is an assumption. Unlike an interlock's failure rate it describes organisational response as much as system capability, and no basis for the specific value is published.
5. Query-time claims are unqualified. Sub-minute query performance is stated without the estate size, retention period or query complexity it was measured against.
6. No cost data. Licence, integration and controller-adapter costs are not reported, so no return-on-investment conclusion is drawn.

> The third limitation is the honest frame for the whole paper. This is an evidence system, and evidence changes outcomes only where someone acts on it.

8 Future Work

- Separating detection from reaction. Instrumenting how often an alert produces a documented intervention would let the 60% modelled reduction be decomposed into system capability and organisational response.
- Query performance envelope. Publishing query time against estate size, retention period and query type would convert a claim into a specification a buyer can size against.
- Predictive drift. The Z-statistic of Equation (drift) detects a mean shift after it occurs; modelling the drift trajectory would allow a station to be flagged before the shift completes.
- Cross-plant capability benchmarking. Aggregating capability indices for equivalent joints across sites would show how much observed variation is process and how much is tooling.
- EPCIS interchange. The fastening record follows the ISO/IEC 19987 event shape but is not yet emitted as a conformant document for exchange with an OEM customer.

9 Conclusion

The facts needed to defend a critical-fastener claim are all measured at the moment the joint is made. What conventional installations lack is a record in which they appear together, bound to the vehicle being built. Assembling them afterwards is a forensic exercise that a warranty deadline does not accommodate, and while it runs the affected population cannot be bounded — so containment is conservative and the over-containment costs more than the defect.

This paper has described a system built so the answer is stored rather than reconstructed: a normalising acquisition layer across five controller vendors, vehicle identity bound at capture through barcode, RFID or PLC handshake, evaluation against the governed joint specification at the moment of fastening, and capability reported per joint with the control state that qualifies it.

The capability reference framework of Section 6.3 is offered as the durable contribution, and its eighth question settles whether a system holds evidence or merely holds data: regenerate a buyoff certificate from six months ago, and see whether it comes back identical.

Appendix A Nomenclature

Table 4. Symbols and abbreviations used in this paper.

Symbol / term	Meaning
USL, LSL	Upper and lower torque specification limits for a joint
mu	Observed process mean of the fastening result distribution
sigma-within	Within-subgroup standard deviation, estimating short-term variation
sigma-overall	Overall standard deviation, including drift between subgroups
Cp, Cpk	Capability indices, valid where the process is in statistical control
Pp, Ppk	Performance indices, used where the process is not in statistical control
Z(k)	Standardised shift of the subgroup mean at subgroup k
n	Subgroup size
A(e)	Admissibility indicator for fastening event e
VIN(e)	Vehicle identification number bound to event e
j(e)	Joint identity of event e; J is the governed joint set
v(e)	Specification version in force when event e was captured
tau(cal)	Time since last calibration of the tool that produced event e
B(VIN)	Backward traceability set — all events bound to a vehicle
F(c)	Forward traceability set — all vehicles affected by condition c
MES	Manufacturing execution system
SPC	Statistical process control

Appendix B Worked Numerical Examples

Appendix B.1 When Cpk and Ppk disagree

A critical joint is specified at 52 Nm with a window of plus or minus 5 Nm, so LSL = 47 and USL = 57. Across one shift, 25 subgroups of four fastenings give a grand mean of 52.4 Nm, a within-subgroup standard deviation of 0.92 Nm, and an overall standard deviation of 1.48 Nm.

Applying Equation (cpk): Cp = (57 — 47) / (6 x 0.92) = 10 / 5.52 = 1.81. The upper margin is (57 — 52.4) / (3 x 0.92) = 4.6 / 2.76 = 1.67; the lower is (52.4 — 47) / (3 x 0.92) = 5.4 / 2.76 = 1.96. Cpk is the smaller, 1.67 — comfortably above any conventional acceptance threshold.

Applying Equation (ppk) with the overall standard deviation: Pp = 10 / (6 x 1.48) = 1.13, and Ppk = 4.6 / (3 x 1.48) = 1.04. The ratio Cpk / Ppk is 1.67 / 1.04 = 1.61, well above 1.

That ratio is the finding. Short-term variation is small, so the station looks excellent on Cpk. Long-term variation is 61% larger, which means the mean is drifting between subgroups — the signature of tool warm-up, a loosening reaction arm, or a joint whose surface condition changes through the shift. Reporting Cpk alone would have recorded a healthy station; ISO 22514-2 directs that Ppk be used here, and Ppk at 1.04 is marginal.

> Same data, same shift, two defensible-looking numbers 60% apart. The control state is not a footnote to a capability report — it determines which number is the report.

Appendix B.2 Catching the drift before it produces a reject

Take the same joint with an established process mean of 52.0 Nm, sigma-within of 0.92 Nm and subgroups of four, so the standard error of the subgroup mean is 0.92 / sqrt(4) = 0.46 Nm.

Successive subgroup means read 52.1, 52.4, 52.8, 53.2, 53.5 Nm. Applying Equation (drift): Z values are 0.22, 0.87, 1.74, 2.61 and 3.26.

Every individual fastening in those subgroups is still inside the 47 to 57 window — the last subgroup sits 3.5 Nm from the upper limit with a spread under 1 Nm, so no result has failed and no operator has seen an alert. But the mean has moved 3.3 standard errors in five subgroups, and continuing at that rate it reaches the specification limit within the shift.

This is the distinction between a maintenance action and a containment action. Flagging at Z = 2.0 produces a tool check on a station still making good joints. Waiting for a reject produces an investigation, a quarantine, and a question about every vehicle built since the drift began.

Appendix B.3 What the forward set is worth

A controller is found to have drifted, and the drift is traced to a maintenance event four days earlier. The line builds 690 units per shift across two shifts, so roughly 5,520 vehicles passed through in that window.

Applying Equation (query) for the forward set: the condition is fastenings performed by that controller identity between the maintenance timestamp and the detection timestamp. Because each event carries a resolved VIN, the query returns the exact list — suppose 1,840 vehicles, since the controller serves one of three parallel stations.

Without identity binding, the containment population is every vehicle built in the window, or 5,520. Binding reduces it to 1,840 — a factor of three — and the 3,680 excluded vehicles are excluded provably rather than by argument. Neither the detection nor the drift changed; only the ability to bound the population did.

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
IATF 16949 audit preparation, attributed to the traceability spine	weeks to one day	**Product specification** /products/torque-traceability — Stated in the Torque Traceability blog post as a consequence of the audit trail generating itself.
Response time to an OEM containment request	minutes, not days	**Product specification** /products/torque-traceability
Forward and backward traceability query time	sub-minute; full per-VIN torque history within an hour	**Product specification** /products/torque-traceability
Fastening errors avoided per year	7,200 of 12,000 (modelled)	**Modelled estimate** 200,000 torqued joints per month, 2.4 million per year; Baseline torque error and rework rate of 0.5%; 60% reduction attributed to the module; Model and defaults published in src/data/roiModels.ts
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.
IATF 16949 audit preparation time	1 day (from 3 weeks)	**Operator-reported — Tier-1 Manufacturers (Pune + Nasik plants, India)** /case-studies/industry40

Metric	Reported value	Provenance
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.

References

[1] 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).

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

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

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

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

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

[7] International Organization for Standardization (2023). ISO 5725-1: Accuracy (trueness and precision) of measurement methods and results - Part 1: General principles and definitions. ISO, Geneva (superseding ISO 5725-1:1994).

[8] International Organization for Standardization (2019). ISO 5725-2: Accuracy (trueness and precision) of measurement methods and results - Part 2: Basic method for the determination of repeatability and reproducibility. ISO, Geneva.

[9] Automotive Industry Action Group (Chrysler, Ford, General Motors) (2010). Measurement Systems Analysis (MSA) Reference Manual. AIAG, 4th edition, MSA-4.

[10] 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).

[11] 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).

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

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

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

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

[16] International Organization for Standardization (2014). ISO 22400-2: Automation systems and integration - Key performance indicators (KPIs) for manufacturing operations management - Part 2: Definitions and descriptions. ISO, Geneva; defines OEE as availability x effectiveness x quality ratio. https://www.iso.org/standard/54497.html

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

[18] International Organization for Standardization (2018). ISO 26262-7: Road vehicles - Functional safety - Part 7: Production, operation, service and decommissioning. ISO, Geneva.

[19] International Organization for Standardization (2015). ISO 13849-1: Safety of machinery - Safety-related parts of control systems - Part 1: General principles for design. ISO, Geneva; a 2023 edition has since been published. https://www.iso.org/standard/69883.html

[20] 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.