Two Clocks, Not One: Escalation Timing, Time-State Capture, and What a Stack Light Cannot Do

Architecture, Computational Methods, and Field Evidence from the MileSoft Andon System

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/andon

Abstract

Andon is usually installed as a light. A stack light above a station turns amber, someone notices, someone comes. The product's own record puts the value of that arrangement at approximately 5% of the methodology — the remaining 95% living in live downtime data, escalation ladders and the analytics they make possible. A light announces a problem; it does not record when the problem started, who came, how long they took, or why it happened, and those four facts are what turn recurring downtime into a fixable one.

This paper presents the architecture and computational methods of the MileSoft Andon System, built on two propositions. First, that the interval between a problem occurring and a responder arriving is a separate quantity from the interval the repair takes, and collapsing them into a single mean time to repair conceals which of the two a plant should invest in. Second, that overall equipment effectiveness is only comparable between plants if the time state is captured on the stop event rather than inferred afterwards from a shift calendar.

We state the escalation state machine and its two independently measured intervals, the ISO 22400-2 decomposition of overall equipment effectiveness with the availability denominator its definition requires, and the Pareto structure that converts an event log into a maintenance programme. The work is anchored in ISO 22400-1 and 22400-2 for the key performance indicators, IEC 62264 for the operations-level placement, and IEC 62541 for programmable-logic-controller integration.

The product's published record reports a 60-80% reduction in mean time to respond and alert latency under 15 seconds with PLC integration. A four-module deployment at a Tier-1 automotive supplier, in which this module was one of four, reported andon mean time to first response falling from six to nine minutes to under fifteen seconds and an eleven percentage-point gain in line-average overall equipment effectiveness; those figures belong to the programme and are not attributed to this module alone.

Keywords: Andon system software; Mean time to first response; MTTR-FR; Escalation ladder; OEE calculation; ISO 22400-2; Downtime tracking; PLC integration OPC UA MQTT; Reason code analysis; Visual factory management; Planned production time; Pareto downtime analysis

1 Introduction

Andon began as a lamp. In its original form on a Toyota line it was a cord an operator pulled and a light that came on, and its genius was social rather than technical: it made a problem visible and gave permission to stop. Most industrial andon installations are still that lamp, sometimes with a buzzer.

The product's own record is blunt about what that delivers: a stack-light andon provides about 5% of the methodology's value, and the other 95% lives in live downtime data, escalation ladders and OEE analytics. The lamp announces; it does not record. It cannot say when the problem started, who responded, how long they took, or what caused it — and without those four facts, a recurring problem stays recurring.

This paper describes a system built for the 95%. Its two organising decisions are that the response interval and the repair interval are measured separately, and that the time state governing the availability calculation is captured on the event rather than inferred from a calendar afterwards.

1.1 Why one mean time to repair is the wrong statistic

A plant reporting a single mean time to repair of eight minutes knows almost nothing actionable. Those eight minutes might be seven minutes of nobody noticing followed by a one-minute fix, or thirty seconds of response followed by seven and a half minutes of genuine repair work. The two situations call for opposite interventions — one is a staffing and notification problem, the other a spares and skills problem.

Separating them is not a refinement of the metric; it is the difference between a number and a diagnosis. The published deployment figures make the point: mean time to first response falling from six-to-nine minutes to under fifteen seconds is an organisational change, and it would be invisible inside a combined figure that also contained repair time.

> A response interval measured from a system-generated raise timestamp cannot be argued with. One reconstructed at shift end is biased toward the incidents nobody was too busy to record — which are the good ones.

1.2 Contributions

1. An escalation state machine with two independently measured intervals, both from system-generated timestamps, so response and repair can be improved separately.
2. Time-state capture on every stop event, making the ISO 22400-2 availability denominator correct by construction rather than by reconstruction.
3. A reason-coded event structure that converts a downtime log into a Pareto ranking traversable to a specific shift, line and cause.

4. An eight-dimension capability reference framework distinguishing an andon system from an alerting device.

2 Background and Related Work

Two bodies of work bear on the design: the standardised definition of the performance indicators an andon system feeds, and the integration reference model that determines where it sits.

2.1 What ISO 22400 fixes, and why it matters here

Overall equipment effectiveness is the most widely reported and least consistently computed number in manufacturing. ISO 22400-2 removes the ambiguity by specifying key performance indicators with their formulae, constituent elements, time behaviour and dimensions. Part 1 supplies the surrounding terminology and the criteria for constructing an indicator at all.

The definition that constrains an andon system is availability: operating time over planned production time. The denominator is not calendar time and not shift length — it is the time the equipment was supposed to be producing. A break, or an interval with no demand, is not downtime; it is outside the denominator entirely.

The consequence is architectural rather than arithmetic. Applying the correct formula to a wrongly bounded denominator produces a number that is standards-shaped but not standards-conformant, and the error is invisible in the result. The time state has to be recorded when the stop happens.

2.2 Where the andon layer sits

IEC 62264, published in parallel as ANSI/ISA-95, places manufacturing operations management at Level 3 — above the control systems of Levels 0 to 2 and below business planning at Level 4. An andon system is unambiguously a Level 3 concern: it consumes control-level signals and produces operations-level information.

Machine-raised alerts arrive over IEC 62541 — OPC Unified Architecture — or over MQTT, from Siemens, Allen-Bradley and Mitsubishi controllers, alongside machine input-output and MES signals. The published record attributes sub-fifteen-second alert latency specifically to this integration path rather than to the escalation logic above it.

> Latency is where machine-raised andon separates from operator-raised andon. A person notices a fault in seconds to minutes; a controller reports it in milliseconds, and the remaining latency is transport and routing.

2.3 Failure modes of light-only andon

- No raise timestamp. Without a system-generated start, the response interval is reconstructed and biased toward incidents that were not urgent enough to distract from recording them.

- One combined interval. Response and repair are collapsed into a single mean time to repair, so a plant cannot tell whether to change staffing or stock spares.
- Inferred time state. Availability is computed against a shift calendar, so planned breaks and no-demand intervals are charged as downtime whenever the calendar and the day disagree.
- Uncoded reasons. Events are recorded without a reason code, producing a downtime total with no Pareto structure and therefore no maintenance programme.
- Escalation by convention. Who responds and when is a matter of local practice, so the ladder exists in people's heads and cannot be measured or changed deliberately.

3 System Overview

The system has three layers: an intake layer accepting operator, machine and sensor triggers; an escalation state machine that owns the two clocks; and a reporting layer producing live displays and the analytics the event log supports.

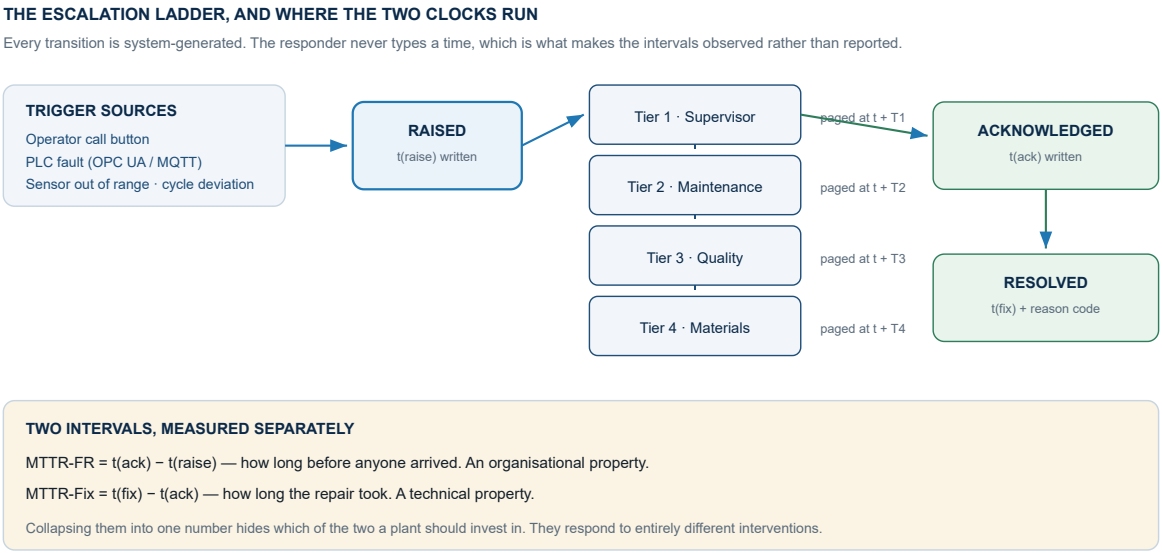

Figure 1. *Schematic.* The escalation ladder and the two clocks. Operator, machine and sensor triggers produce the same raised state, whose system-generated timestamp starts the response clock. Successive tiers are paged automatically as thresholds elapse. Mean time to first response ends at acknowledgement; mean time to fix runs from acknowledgement to resolution. Neither timestamp is entered by a responder.

3.1 Intake

Three trigger classes produce identical events. Operators flag quality, material or machine issues via call buttons. Machines raise alerts automatically through programmable-logic-controller fault signals or cycle-time deviation. Sensors trigger on out-of-specification readings.

Normalising all three into one event stream is what makes the analytics coherent. A plant whose operator calls and machine faults live in separate systems has two downtime records that cannot be ranked against each other, and the largest single cause is frequently the one split across both.

3.2 The escalation state machine

An event moves through raised, acknowledged and resolved. Escalation tiers — supervisor, maintenance, quality, material handling — are paged automatically as configured thresholds elapse from the raise timestamp, so an unacknowledged event climbs the ladder without anyone deciding to escalate it.

The response clock starts at raise rather than at acknowledgement, which is the detail that makes the measurement honest. Starting it at acknowledgement would measure how long the responder took after noticing, and would report an excellent figure for an event nobody noticed for six minutes.

3.3 Display and record

Large overhead LED displays, line-side monitors, supervisor desktop dashboards and mobile notifications all render the same real-time event stream. The display layer is deliberately thin: it presents state rather than holding it, so a display failure degrades visibility without losing the record.

Every event is logged with reason code, line, station, shift and resolution time. That five-tuple is what the analytics of Section 4.3 traverse, and it is why the published record can describe answering a question about one shift on one line in three clicks.

4 Computational Methods

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

4.1 The two intervals

Both intervals are differences between system-generated timestamps. Neither is entered by a person, and neither can be adjusted after the fact.

$$\mathit{MTTR\text{-}FR} = \mathit{mean\ of}\ [\ t(\mathit{ack}) - t(\mathit{raise})\ ],\ \mathit{MTTR\text{-}Fix} = \mathit{mean\ of}\ [\ t(\mathit{fix}) - t(\mathit{ack})\ ] \tag{mttrfr}$$

t(raise) is when the event was written to the store, t(ack) when a responder acknowledged it, and t(fix) when the resolution and its reason code were recorded. Their sum is the conventional single mean time to repair, which this system reports as two numbers rather than one.

Reporting both changes what a target means. A plant can set a response-time target that its supervisors control and a repair-time target that its maintenance function controls, and each is answerable for a quantity it can actually move.

4.2 Overall equipment effectiveness, with the correct denominator

Overall equipment effectiveness decomposes into three ratios, each isolating a different loss class, exactly as ISO 22400-2 defines them.

$$OEE = A \times E \times Q, \text{ where } A = \text{operating time} / \text{planned production time}, E = (\text{reference cycle time} \times \text{produced quantity}) / \text{operating time}, Q = \text{good quantity} / \text{produced quantity} \quad \text{(oee)}$$

planned production time excludes intervals the equipment was not scheduled to run. Operating time is planned production time less the stops that occurred within it.

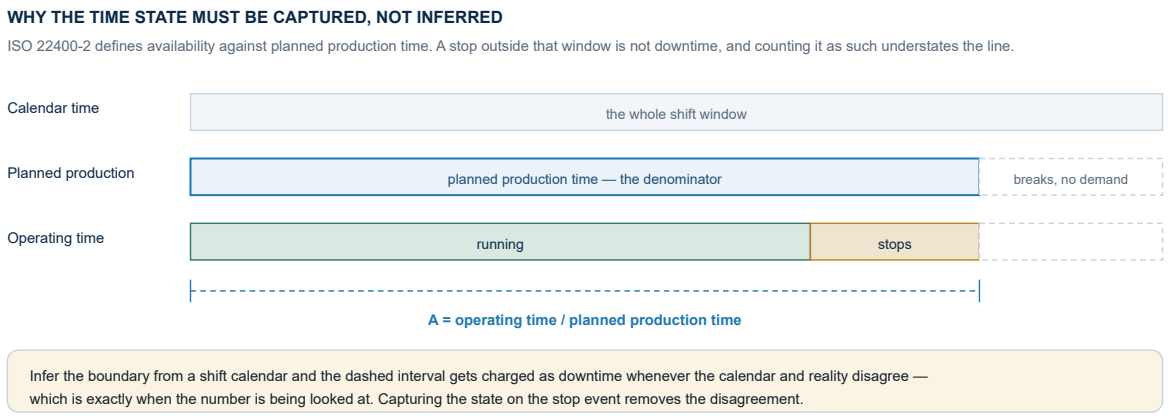

Figure 2. *Schematic.* The time-state model. Planned production time is bounded by what the equipment was scheduled to do, not by the shift. When the boundary is inferred from a calendar rather than captured on the event, the dashed interval is charged as downtime whenever calendar and reality disagree — which is exactly when someone is examining the number.

> This is the difference between an OEE figure that is comparable between two plants and one that is merely computed with the right formula.

4.3 From event log to maintenance programme

Downtime cost concentrates. A small number of reason codes typically account for most lost minutes, and the analytics rank them so that effort follows the concentration rather than the complaint volume.

$$D(\text{reason } r) = \text{total lost minutes across all events with that reason code}, \pi(r) = \text{that reason's share of all lost minutes} \quad \text{(pareto)}$$

the total interval from raise to fix is used rather than the repair interval alone, because the line was stopped for both. Ranking by pi(r) rather than by event count is what prevents a frequent but trivial cause from outranking a rare but expensive one.

The same structure supports the traversal the published record describes — a question about one shift on one line answered in three clicks — because reason code, line, station and shift are dimensions of one record rather than fields spread across systems.

4.4 Where the latency budget goes

The published sub-fifteen-second alert latency is a budget across four stages, and it is worth decomposing because only one of them is under the andon system's control.

$$L(\text{alert}) = \text{detection at the controller} + \text{transport over OPC UA or MQTT} + \text{routing through the escalation rules} + \text{notification delivery} \quad \text{(latency)}$$

detection and transport are properties of the control network; routing is the andon system's own; notification depends on the delivery channel, and a mobile push is not equivalent to a line-side display in either latency or reliability.

Stating the budget this way makes an important limitation visible: a system quoting an end-to-end latency figure is quoting a number that depends heavily on a network it does not own. The comparable figure between vendors is the routing term.

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

The product's own record reports a 60-80% reduction in mean time to respond once andon is integrated with the plant's response organisation, and alert latency under fifteen seconds with programmable-logic-controller integration over OPC UA or MQTT.

The qualifier in the first claim carries weight. The reduction is attributed to andon plus an integrated response organisation, not to the software alone — which is consistent with the architecture, since an escalation ladder that pages a tier nobody staffs shortens no interval at all.

The record's most useful statement is not a percentage. It is the assertion that a stack-light andon delivers about 5% of the methodology's value, with the remaining 95% in live downtime data, escalation ladders and OEE analytics. That is a claim about where value sits rather than about the product, and the whole architecture follows from it.

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 the Andon System was installed alongside Torque Traceability, Material Traceability and the manufacturing ERP. It replaced a physical stack light with a paper escalation sheet. The programme reported andon mean time to first response falling from six-to-nine minutes to under fifteen seconds, and an eleven percentage-point gain in line-average overall equipment effectiveness.

> 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 response-time figure is the one most plausibly attributable here, since it measures an interval this module owns end to end and the case study names the prior mechanism it replaced. The eleven-point OEE gain is much less separable: availability, effectiveness and quality ratio are all affected by the other three modules.

One methodological caution applies to the response figure as it does to the dwell-time figures in the warehouse paper. The six-to-nine minute baseline came from a paper escalation sheet; the post-deployment figure is system-generated. Some of the improvement is a measurement change rather than an operational one, and the published data does not separate the two.

5.3 Modelled downtime recovery

Table 1. Published model parameters and the default-case result. Every input is adjustable; the result is a modelled estimate rather than a measurement.

Parameter	Default
Unplanned downtime per month	40 hours
Annual unplanned downtime	480 hours
Reduction attributed to the module	70%
Downtime hours eliminated per year	336

The 70% assumption sits at the upper end of the product's own published 60-80% range, and the model applies it to total unplanned downtime rather than to the response interval alone. That is a broader claim than the published range supports: shortening response time reduces the raise-to-fix interval, but it does not shorten the repair itself. The model is therefore the least conservative of those in this series and should be treated as an upper bound.

6 Discussion

6.1 Andon is an organisational instrument with a software body

The response interval is not primarily a technical quantity. Once alert latency is under fifteen seconds, essentially all of the remaining response time is people deciding to move. A plant that deploys this system and staffs no tier will measure its own inaction precisely and change nothing.

This is why the product's own record qualifies its headline figure with the phrase about an integrated response organisation, and it is worth repeating rather than eliding. The software's contribution is to make the interval visible, attributable and comparable across shifts — which is a precondition for improvement rather than improvement itself.

6.2 The denominator is where OEE comparisons go wrong

Two plants can apply the ISO 22400-2 formula faithfully and report OEE figures that are not comparable, because one bounded planned production time by a shift calendar and the other by what was actually scheduled. The formula is not where the divergence enters.

The correction is unglamorous: record the time state on the stop event. It costs a field and a decision at capture, and it is the difference between a number a plant can benchmark against another site and a number that only compares against its own history — and only while its calendar stays accurate.

6.3 A capability reference framework for andon systems

Table 2. Capability reference framework. Each dimension distinguishes an andon system from an alerting device, and each is answerable by demonstration on a running line.

Dimension	Question the system must answer by demonstration
D1 Two clocks	Are response and repair reported as separate intervals, or as one combined figure?
D2 Generated timestamps	Is any of raise, acknowledge or resolve entered by a person?
D3 Clock start	Does the response interval start at raise or at acknowledgement?
D4 Time state on the event	Does a stop record whether it fell inside planned production time, or is that inferred later?
D5 Unified intake	Do operator calls and machine faults produce the same event type in one stream?
D6 Automatic escalation	Leave an event unacknowledged. Does the next tier page without anyone deciding?
D7 Reason-coded Pareto	Rank causes by lost minutes, not event count. Can you reach one shift on one line in three steps?
D8 Latency decomposition	Is the quoted alert latency broken into detection, transport, routing and notification?

D3 is the question a vendor is least likely to volunteer. A response clock starting at acknowledgement reports how fast someone worked once they noticed, which is not the interval a plant is trying to shorten.

6.4 Generalisability

The two-clock structure generalises to any incident-response setting where arrival and remedy are distinct — maintenance, laboratory operations, IT service management. The time-state argument is specific to OEE and to equipment with a schedule; it does not apply where equipment runs continuously. The 60-80% response reduction was reported in a plant replacing a paper escalation sheet, and a plant already running a digital escalation ladder should expect far less.

7 Threats to Validity and Limitations

1. Vendor-reported figures. The 60-80% response reduction and the sub-fifteen-second latency are published by the vendor without sample size, distribution or independent measurement.
2. Baseline measurement asymmetry. The deployment's six-to-nine minute baseline came from a paper escalation sheet while the post-deployment figure is system-generated, so part of the improvement is a measurement change. The published data does not separate the two.
3. Confounded deployment result. The eleven percentage-point OEE gain belongs to a four-module programme; availability, effectiveness and quality ratio are all affected by the other three modules.
4. The modelled 70% is an upper bound. It applies the top of the published response-reduction range to total unplanned downtime, whereas shortening response does not shorten repair.
5. Latency depends on a network the system does not own. Detection and transport dominate the budget in Equation (latency) and are properties of the control network.
6. The 5% figure is a characterisation, not a measurement. That a stack light delivers about 5% of andon's value is a published judgement about where value sits, offered without derivation.
7. No cost data. Display hardware, PLC integration and licence costs are not reported, so no return-on-investment conclusion is drawn.

> The second limitation applies to most before-and-after response-time claims in this industry, including this one. A plant migrating from paper to system-generated timestamps should expect its first honest number to look worse than its last dishonest one.

8 Future Work

- Parallel-capture baselines. Running paper and system-generated timestamps together for one period would bound the measurement artefact described in Section 7 and let response improvements be stated net of it.
- Routing latency as a published specification. Decomposing the alert budget and publishing only the routing term would give buyers a figure that is actually comparable between vendors.
- Response-time distribution rather than mean. A mean conceals the tail, and the tail is where the expensive incidents sit; reporting a percentile would target the events that matter.
- Escalation-threshold learning. Tier thresholds are configured per site; deriving them from observed acknowledgement distributions would set them where they actually change behaviour.
- Isolating availability losses in multi-module deployments, so an andon contribution can be separated from concurrent quality and traceability work.

9 Conclusion

A stack light announces a problem and records nothing about it. The product's own record puts that at roughly 5% of what the andon methodology offers, and this paper has taken that judgement seriously enough to build an architecture around the other 95%.

Two decisions carry it. Response and repair are measured as separate intervals from system-generated timestamps, so a plant can tell whether its problem is that nobody comes or that the fix is slow — which call for opposite investments. And the time state is captured on every stop event rather than inferred from a calendar, which is what makes an availability figure conformant with ISO 22400-2 rather than merely formula-shaped.

The capability reference framework of Section 6.3 is offered as the durable contribution, and its third question is the one a buyer should ask first: does the response clock start when the event was raised, or when somebody noticed it?

Appendix A Nomenclature

Table 3. Symbols and abbreviations used in this paper.

Symbol / term	Meaning
t(raise)	System timestamp when the event was written to the store
t(ack)	System timestamp when a responder acknowledged the event
t(fix)	System timestamp when resolution and reason code were recorded
MTTR-FR	Mean time to first response — raise to acknowledgement
MTTR-Fix	Mean time to fix — acknowledgement to resolution
T(ppt)	Planned production time — the ISO 22400-2 availability denominator
T(op)	Operating time — planned production time less stops within it
A, E, Q	Availability, effectiveness and quality ratio
OEE	Overall equipment effectiveness, the product A x E x Q
t(ref)	Reference cycle time per unit
q(prod), q(good)	Produced and good quantity in the period
D(r)	Total lost minutes attributed to reason code r
pi(r)	Reason code r's share of all lost minutes
L(alert)	End-to-end alert latency, decomposed in Equation (latency)
PLC	Programmable logic controller
MQTT	A lightweight publish-subscribe transport used for machine telemetry

Appendix B Worked Numerical Examples

Appendix B.1 Why one number hides the answer

Two lines each report a combined mean time to repair of 8.0 minutes across a month. Line A records mean time to first response of 7.2 minutes and mean time to fix of 0.8 minutes. Line B records 0.5 and 7.5 minutes.

Applying Equation (mttrfr) to each gives the two pairs above; their sums are identical at 8.0 minutes, which is the only figure a conventional report would show.

The interventions are opposite. Line A has a notification and staffing problem: the fix takes under a minute once someone arrives, so escalation thresholds, tier coverage and display placement are where the time is. Line B has a technical problem: response is essentially instant and the repair is slow, pointing at spares availability, tooling or skills.

> A plant that invested in faster notification on Line B would spend money to improve 0.5 minutes out of 8.0. The combined figure gives no way to know that.

Appendix B.2 What the denominator does to the number

A shift runs 480 minutes. Scheduled breaks account for 30 minutes and a period of no demand for a further 45, so planned production time is 480 — 75 = 405 minutes. Recorded stops within planned production time total 58 minutes, so operating time is 405 — 58 = 347 minutes. Reference cycle time is 0.55 minutes per unit, produced quantity 600, good quantity 591.

Applying Equation (oee): A = 347 / 405 = 0.857. E = (0.55 x 600) / 347 = 330 / 347 = 0.951. Q = 591 / 600 = 0.985. OEE = 0.857 x 0.951 x 0.985 = 0.803, or 80.3%.

Now infer the denominator from the shift calendar instead, so all 480 minutes count as planned and the 75 minutes of breaks and no demand are charged as downtime alongside the 58 minutes of real stops. Operating time is unchanged at 347 minutes, but A = 347 / 480 = 0.723 and OEE falls to 0.723 x 0.951 x 0.985 = 0.677, or 67.7%.

A 12.6 percentage-point difference, produced entirely by how the denominator was bounded. Neither the line nor the operators changed. A plant benchmarking the second figure against a peer reporting the first would conclude it had a serious availability problem it does not have.

Appendix B.3 Ranking by minutes, not by frequency

A month's log holds four reason codes. Material shortage: 61 events, 214 minutes total. Tool change overrun: 12 events, 396 minutes. Quality hold: 9 events, 288 minutes. Conveyor jam: 44 events, 132 minutes.

By event count the ranking is material shortage, conveyor jam, tool change, quality hold. Applying Equation (pareto) by lost minutes the total is 1,030 and the shares are: tool change 38.4%, quality hold 28.0%, material shortage 20.8%, conveyor jam 12.8%.

The rankings are nearly inverted. Material shortage generates five times as many events as tool change overrun and costs roughly half as many minutes. A programme driven by complaint volume would work on the wrong cause first — which is the specific failure a reason-coded log with durations prevents.

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
Mean time to repair reduction attributed to Andon	60-80%	**Product specification** /products/andon
Alert latency with PLC integration over OPC UA or MQTT	under 15 seconds	**Product specification** /products/andon
Share of the andon methodology's value delivered by a stack light alone	approx. 5%	**Product specification** /products/andon — The remaining 95% is attributed to live downtime data, escalation ladders and OEE analytics.
Unplanned downtime hours eliminated per year	336 of 480 h (modelled)	**Modelled estimate** 40 downtime hours per month, 480 per year; 70% reduction in unplanned downtime and MTTR attributed to the module; Model and defaults published in src/data/roiModels.ts
Overall equipment effectiveness, line average	+11 percentage points	**Operator-reported — Tier-1 Manufacturers (Pune + Nasik plants, India)** /case-studies/industry40
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.
Phased rollout duration	8 weeks (plant 1), 6 weeks (plant 2)	**Operator-reported — Tier-1 Manufacturers (Pune + Nasik plants, India)** /case-studies/industry40

References

[1] 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

[2] 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

[3] 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).

[4] 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).

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

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

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

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

[10] 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

[11] 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

[12] 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

[13] 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

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

[15] 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.

[16] 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).

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

[18] 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

[19] 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.

[20] Shingo, S. (translated by A. P. Dillon) (1986). Zero Quality Control: Source Inspection and the Poka-Yoke System. Productivity Press, Cambridge, MA. Originally published in Japanese, 1985.