




LEVARU WHITE PAPER

The Closure Note

A Facility Manager's Framework for Capturing the Diagnosis, the Fix and the Exception Before the Technician Leaves the Room

Nikolai Hanov

Published by Levaru · May 2026



LEVARU

levaru.co · Alexandria, Virginia · +1 (703) 646-8300

RESEARCH SERIES · PAPER 1

Abstract

Facility operations produce knowledge continuously and record it almost never. Every completed work order ends at the same moment: the fault is understood, the fix has held, and the one person who knows both is standing at the asset with a text field open — a field most organizations allow to close with a status and a single word. This paper defines the closure note as an operational record and closure discipline as the practice that produces it: a five-line anatomy capturing what was found, what caused it, what was done, what was learned and what comes next; capture mechanisms matched to the technicians who must actually use them; a sampled Closure-Quality Audit scored on the Recurrence-Recoverability Matrix; and the Re-Fix Test as the headline measure of whether the archive can answer the building's next question without a phone call. It locates the discipline's first paybacks at three moments — the recurring fault, the handoff and the arrival of the external reader — and supplies a staged implementation path, a five-level maturity model, a baseline-first metrics discipline, an explicitly illustrative cost model that counts the discipline's three-minute tax honestly, and the governance that keeps a record about assets from ever becoming a dossier on people. It aligns the practice with ISO 41001, the ISO 55000 series, ISO 14224's failure-coding logic, EN 15341 and SMRP metrics conventions, COBie handover thinking and IFMA's competency framework, and closes with candid limitations.

KEYWORDS

work orders closure notes CMMS data quality failure coding maintenance records
knowledge capture onboarding ISO 14224 ISO 41001 facility management

Suggested citation: Hanov, N. (2026). *The Closure Note: A Facility Manager's Framework for Capturing the Diagnosis, the Fix and the Exception Before the Technician Leaves the Room*. Levaru. levaru.co/white-papers/closure-note/

About this series: This paper is part of Levaru's research series on the operational record of buildings, alongside *The Visual Asset Register*, *Minimum Viable Digital Twin for Existing Buildings* and *The Continuity Audit*.

Vendor neutrality: This paper names no commercial platform and describes capability categories rather than implementations. © 2026 Levaru. All rights reserved.

Contents

1. The problem: work completed, knowledge discarded	4
2. Definition: what a closure note is — and is not.....	5
3. Why now: three converging pressures	6
4. What makes it a discipline, not a habit: five properties	7
5. The instruments.....	8
5.1 The Re-Fix Test.....	8
5.2 The five-line anatomy	9
5.3 The Closure-Quality Audit and the Recurrence-Recoverability Matrix	10
5.4 The capture-mechanisms menu	11
6. Three pillars: where the record pays back first.....	12
6.1 The recurring fault	12
6.2 The handoff and the new hire.....	12
6.3 The contract, the warranty and the audit.....	13
7. Implementation: crawl, walk, run.....	13
8. A maturity model for the work-order record	15
9. Measuring the program: baselines before benchmarks.....	16
10. Governance: a record about assets, not a dossier on people.....	18
11. Standards and frameworks alignment.....	19
12. Limitations and outlook.....	20
13. Conclusion	20
References	21

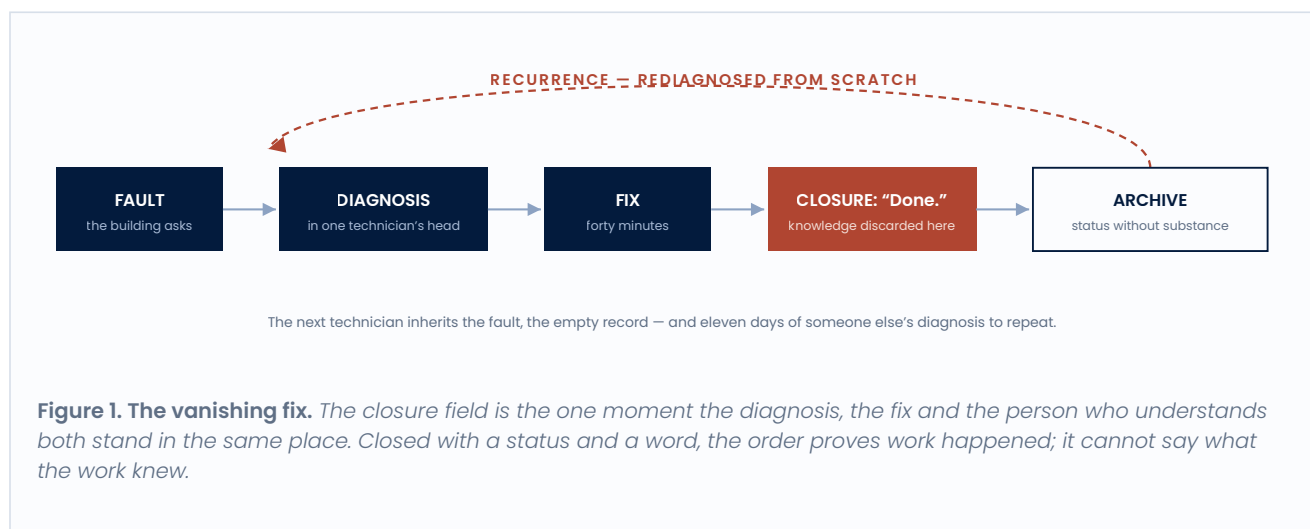
1. The problem: work completed, knowledge discarded

At 6:40 on a February evening, a technician closes out the fault that has intermittently tripped the penthouse air handler for three weeks. The diagnosis was earned the hard way: two false starts chasing the freeze-stat, a morning spent watching the unit through its warm-up cycle, and finally the observation that the outside-air damper linkage was binding at a third of its travel — starving the coil just long enough, on cold starts, for the safety to catch. The fix took forty minutes. The understanding behind it took eleven days. The work order that records both closes with a status of *Completed* and a note that reads, in full: “Reset unit. Done.”

Nothing in that transaction involves anyone failing at their job. The technician was hired to restore the equipment, the equipment is restored, and the system of record has been fed exactly what it demands — a status, a timestamp, a labor line. But something of real and measurable value has just been discarded at the moment of its creation. The diagnosis, the cause, the trick of the fix and the warning for the next cold snap existed, complete and fresh, in exactly one place: the mind of a person standing at the asset with a text field open. The field closed around one word, and the knowledge walked to the parking lot.

“The most expensive sentence in facility management is ‘Done.’”

Multiply the transaction by an operation’s history and the shape of the problem emerges. A mid-sized portfolio’s CMMS holds tens of thousands of completed work orders — statuses immaculate, dates precise, text fields empty or nearly so. It is an archive that can prove work happened without being able to say what the work knew. Levaru’s earlier research called the missing layer *located knowledge*: knowledge that is real, operational and tied to specific places, assets and relationships — and almost never written down, because for as long as its holder is one radio call away, it never needs to be [1]. The work order is where located knowledge surfaces on a schedule — every fault summons it — and the closure field is the one moment in ordinary operations when the problem, the solution and the person who understands both stand in the same place at the same time. Most organizations let the moment pass, every day, at the cost of the next recurrence.



The loss is discovered by three readers, each arriving later than the moment that could have served them. The first is the **next technician**: the fault recurs — because fault families recur, seasonally and mechanically — and the responder opens the asset history, finds three completed orders and no content, and begins the eleven days again. The second is the **successor**: a new hire, a covering colleague, the person who inherits the portfolio after a retirement or a reorganization, for whom the archive should have been a syllabus and is

instead a list of dates. The third is the **external reader**: the incoming contractor at a changeover, the warranty adjuster, the insurer, the auditor — strangers with standing to ask what was done and when, for whom an empty record is at best a lost argument and at worst a finding.

This paper proposes that the remedy is neither a platform nor a project but a discipline, operated at the point where the knowledge already surfaces. It defines the closure note as an operational record with a stated anatomy; supplies the instruments that make note quality visible, comparable and improvable — the Re-Fix Test, the Closure-Quality Audit and the Recurrence-Recoverability Matrix; matches capture mechanisms to the technicians who must use them; and sets out the implementation path, metrics, governance and standards alignment that turn a good intention into an operating practice. The argument throughout is the one this research series has made about the operational record generally [1][2]: the knowledge is not missing — it is present, briefly, on a schedule — and the organization's only real choice is whether to be holding a pen when it appears.

2. Definition: what a closure note is — and is not

Definition. A **closure note** is the completion record of a work order that captures, in a structured minimum, what was found, what caused it, what was done, what was learned and what should happen next — written at the point of work, filed against the asset in the system of record, and sufficient for a competent stranger facing the same fault to reach the same fix without asking anyone. **Closure discipline** is the operating practice — anatomy, capture mechanisms, sampled audits, feedback and governance — that makes such notes the ordinary output of ordinary work rather than the occasional gift of the conscientious.

Each clause of the definition is load-bearing. **A structured minimum** means the note has a stated anatomy — Section 5.2 gives it five lines — and that the anatomy scales with criticality: the failed lamp does not need five lines, and the chiller does, and the difference is a documented judgment rather than a mood. **At the point of work** means the note is written, dictated or photographed while the technician still stands at the asset, because the memory of a repair decays in hours and the batch of orders closed from a desk on Friday afternoon is a work of fiction with accurate timestamps. **Against the asset in the system of record** means the note lands where the next reader will look — attached to the equipment record, searchable, in a system that survives its author. A superb note in a personal notebook, a phone's photo roll or a technician's memory of a text thread fails the definition completely: for continuity purposes it does not exist, a standard this series has applied to every layer of the operational record [1]. **A competent stranger** imports the test that anchors the series' measurement of documentation generally — the Handoff Test of Levaru's earlier work [1] — and applies it at the level of the single fix: the reader to write for is not the author in a month but a qualified outsider with no one to call. And **the ordinary output of ordinary work** is the clause that separates a discipline from a virtue: the practice cannot depend on hiring writers, and does not — it depends on an anatomy small enough to complete, mechanisms matched to the people completing it, and a feedback loop that notices.

The closure note sits deliberately between two established responses, and is defined as much by what it declines from each as by what it takes. At one pole is the **status-only archive** — the CMMS operated as a timesheet, where completion is a checkbox and the text field is decoration. It is cheap, universal and structurally amnesiac: it can reconstruct labor hours and cannot reconstruct a single diagnosis. At the other pole is the **documentation project** — the SOP library, the wiki mandate, the instruction that every team write down its procedures. These are systematic but parallel: they run beside the workflow rather than inside it, they are always deferrable and therefore usually deferred, and they ask for knowledge in the abstract, at a desk, months from the moment it was live. Closure discipline rides the workflow that already exists. The order is

already open; the technician is already at the asset; the knowledge is already in hand. The discipline adds minutes, not meetings — and Section 9 counts those minutes honestly, because they are the practice's real and permanent cost.

Two boundary clarifications complete the definition. First, the closure note records operating reality, not design intent: it is not a substitute for drawings, O&M manuals or engineering documentation, and it will frequently record the places where the building has quietly diverged from all three — which is much of its value. Second, the note is raw material for analysis, not the analysis itself: it feeds root-cause work, reliability trending and capital argument, but a closure discipline is not an RCA program, and pretending otherwise burdens a five-line record with a committee's job. What the note owes the future is evidence — accurate, contemporaneous, findable. What the organization does with the evidence remains, as it should, a decision.

3. Why now: three converging pressures

None of the pressures behind closure discipline is new. Buildings have always presented the same faults twice, and every maintenance culture has told a version of the story about the note that would have saved the night shift. What has changed is the arithmetic: the rate at which the archive's authors are leaving, the measured price of knowledge that must be asked for, and the number of outside readers a work-order archive must now satisfy have all moved in the same direction at once. Three published pressures — each cited from independent sources rather than generated by this paper — frame the moment.

The first pressure is that **the archive's authors are leaving**. A study cited in IFMA's *Facility Management Journal* found that 31 percent of facility management employees are older than 55 and eligible to retire between 2022 and 2037 [8] — nearly a third of the people whose closure notes could have been written over the past decades, on a countdown already several years advanced. The pattern is visible wherever institutions publish their numbers: the University of Tennessee, Knoxville's facilities organization — roughly 600 staff across more than 250 buildings — recorded 97 retirements between 2020 and 2025 [9], as discussed in this series' earlier work [1]. Research on aging workforces has warned since at least DeLong's *Lost Knowledge* that organizations systematically underestimate what leaves with long-tenured staff [10]. An empty archive is tolerable for exactly as long as its would-be authors remain one radio call away; every departure converts a deferred question into a permanent one, and the work-order record is the largest single place the deferral has been accumulating.

The second pressure is **the measured cost of asking**. The Panopto Workplace Knowledge and Productivity Report, a study of more than 1,000 United States workers conducted with YouGov, estimated that 42 percent of institutional knowledge is unique to the individual who holds it, and put the friction cost of that arrangement at 5.3 hours per employee per week spent waiting for information from colleagues or recreating knowledge that already exists somewhere in the organization [3]. The figures are consistent with the broader information-friction literature this series has cited before: McKinsey's estimate that interaction workers spend roughly 19 percent of the workweek searching for and gathering information [4], and a 2023 Quickbase survey in which 53 percent of respondents reported chasing information more than 10 hours a week [5][1]. Facility work adds a physical dimension — the waiting happens on ladders, and the recreating happens through repeat site visits and re-run diagnoses. A team does not need to adopt any of these numbers as its own; Section 9 insists on measuring the team's actual baseline. But their consistency establishes the shape of the problem: every empty closure field converts a one-time writing cost of minutes into a recurring asking cost of hours, payable by whoever meets the fault next.

The third pressure is that **the archive has acquired external readers**. Grand View Research valued the global facility management services market at approximately US\$1.75 trillion in 2024, with outsourced delivery accounting for 61.5 percent of revenue [7] — which means that in a majority of the market, the work-order

archive is the principal artifact one organization hands another at every contract changeover, and the incoming provider's first month is spent reading whatever the outgoing provider's technicians wrote, or did not. Warranty claims, insurance recoveries and disputed handovers stand or fall on contemporaneous records; regulators and auditors increasingly ask not whether maintenance was scheduled but whether its execution can be evidenced. And deferred-maintenance pressure sharpens the stakes for the archive's internal readers too: Gordian's *State of Facilities in Higher Education* puts deferred capital renewal on North American campuses at US\$156 per gross square foot [6], and a renewal case argued from documented failure history reads very differently — to a CFO or a board — from one argued from anecdote.

THE PRESSURE, IN PUBLISHED FIGURES

42%

of institutional knowledge is unique to the individual who holds it — and to no one, once they are gone [3]

19%

of the interaction worker's week is spent searching for and gathering information [4]

53%

of surveyed workers chase information more than 10 hours per week [5]

US\$156

of deferred capital renewal per gross square foot on North American campuses — a case that needs evidence [6]

Sources: Panopto / YouGov [3]; McKinsey Global Institute [4]; Quickbase [5]; Gordian [6]. Full citations in the References.

Three readers, one field. The next technician, the successor and the stranger with a contract will each, in time, open the archive and ask it a question. The organizations most exposed are the ones whose archives can prove that work happened without being able to say what the work knew — and producing that evidence, one order at a time, is the discipline's entire job.

4. What makes it a discipline, not a habit: five properties

Many organizations have gestured at this problem before. There have been “notes required” policies enforced by a mandatory field that accepts a single period; memos from a new supervisor about documentation standards, observed for a quarter; templates designed in an office and abandoned in a plant room; and archives that hold, side by side, the meticulous records of one conscientious technician and the timestamps of everyone else. Closure discipline is distinguished from these gestures by five properties. Remove any one of them and the exercise reverts to a habit — admirable where it occurs, unmeasured, unowned, and gone within two rotations of staff.

Property one: a defined anatomy, not an exhortation. “Write better notes” is a wish; five named lines — found, cause, action, learned, next — are a standard, teachable in an hour and auditable in a glance. The anatomy pairs prose with a small failure-code taxonomy in the manner the reliability world standardized long ago [12]: *codes for counting, prose for knowing*. Codes make ten orders comparable and a chronic asset visible in a trend; prose makes one order recoverable by a stranger at 2 a.m. Each is routinely asked to do the other's job, and fails. And the anatomy scales deliberately: a documented three-line minimum — and, below a stated threshold, a one-line standard — keeps the discipline honest about the difference between a chiller and a lamp, which is precisely the judgment blanket policies cannot make.

Property two: captured at the point of work. The half-life of repair memory is measured in hours. The technician who dictates thirty seconds at the asset produces a better record than the one who types for five minutes at a desk on Friday, because the desk version is a reconstruction and the asset version is a report. The discipline therefore treats capture mechanics as a first-class design problem — Section 5.4 supplies the menu — and treats end-of-day batch closure as what it is: the manufacture of plausible fiction with accurate timestamps.

Property three: filed against the asset, findable by a stranger. A note earns its keep only where the next reader will look: attached to the asset record, searchable by equipment and by fault, in the system of record rather than beside it. Mapping where notes actually land produces the discipline's quietly valuable by-product — the same gap map this series has described for the operational record generally [1]: a candid picture of which systems the team's real knowledge lives in and which are ceremonial. A photographed fix in a private message thread is a favor to one colleague; the same photograph on the asset record is infrastructure.

Property four: sampled, scored and fed back on a cadence. The discipline is kept alive by a small standing audit — Section 5.3 — that samples a few dozen closed orders a month, scores them on two coarse axes, and feeds the findings to exactly three places: the template library, the toolbox talk and the enrichment backlog. Feedback goes to teams and to the standard, never to individual files; Section 10 treats that boundary as non-negotiable. An unaudited note standard decays to the empty field at the speed of staff turnover, which is the speed that made the standard necessary.

Property five: governed handling of a people-adjacent record. Closure notes carry authorship, timestamps and — when the discipline is working — honest accounts of what was found, including what a previous repair got wrong. That content is gathered from and about individuals, and it stays honest only under written rules established before the standard is announced: what the notes will and will not be used for, who sees quality findings at what level of aggregation, and an explicit, kept commitment that nothing in the record feeds performance evaluation or pace surveillance. Governance is listed here as a defining property rather than a compliance layer, because it is a condition of the archive telling the truth.

Closure discipline is	It is not
A structured completion record — five lines, scaled to criticality, written at the asset	A time-and-motion study, a pace log or any instrument of surveillance
Sampled, scored and fed back on a cadence, to teams and templates	A root-cause-analysis program — it feeds one; it is not one
Filed against the asset in a system of record a stranger can search	A documentation project or wiki mandate running beside the workflow
Operated through mechanisms the technicians actually chose	A software purchase — every mechanism on the menu runs on commodity tools
Governed — evaluation firewall, just-culture handling, access control	A promise of literary prose — the standard is <i>useful to a stranger</i> , not eloquent

The discipline at a glance: five defining properties against five stated exclusions.

5. The instruments

Four instruments turn the definition into practice. None requires new software; all four require the honesty of whoever is holding the pen.

5.1 The Re-Fix Test

Levaru's earlier work proposed the Handoff Test as the headline measure of the operational record: what share of a role's located knowledge could a competent stranger recover from the organization's systems alone [1]? This paper supplies its work-order-level sibling. Where the Handoff Test asks whether a *role* survives its holder, the Re-Fix Test asks whether a *fix* survives its fixer.

The Re-Fix Test. If this fault recurred tomorrow on this asset — and the technician who resolved it were unreachable — could a competent stranger, reading only the work-order record, reach the same diagnosis and the same fix without re-deriving either? The share of a sampled set of closed orders that passes, honestly judged, is the team's **Re-Fix score**. The distance between that score and 100 percent is knowledge the organization paid to acquire and declined to keep.

Applied as an instrument rather than a thought experiment, the test acquires a method. Each month, the audit of Section 5.3 pulls a sample of recently closed corrective orders and judges each against the record alone — the note, the codes, the attached photographs, the asset history — with a deliberately binary verdict: a stranger could re-fix from this, or could not. The score can be weighted by criticality, so that a team whose passing notes are all lamp replacements scores as badly as it should. The test can also be run live, and periodically must be: hand a competent outsider — a technician from another building, a new hire past orientation — the fault description from a genuinely closed order, restrict them to the record, and watch where it fails them. Twenty minutes of observed struggle settles arguments that a month of self-assessment cannot, and teams that run the live version rarely dispute their score again.

Three properties make the Re-Fix Test the right headline measure. It is **order-anchored**, so it accumulates one honest verdict at a time and survives every individual departure. It is **directional**, moving only when notes actually improve in the systems a stranger can reach — it cannot be gamed by policy announcements. And it is **bracing**: most teams sampling their archive for the first time produce a score they decline to publish, which is precisely the information the organization did not previously have.

5.2 The five-line anatomy

The anatomy is the discipline's core artifact: five named lines, teachable in a toolbox talk, completable in the minutes the fix already earned. Each line exists because a specific future reader needs it.

FOUND records the symptom and the as-found state — what the responder walked into, and what the diagnosis cost to establish. It is the line that would have saved the eleven days. **CAUSE** records what actually failed, as distinct from what complained: the binding linkage, not the tripping safety. It is the line the failure code attaches to, and the line that turns three isolated events into one chronic asset. **ACTION** records what was done — steps, parts, settings, the as-left state — in enough detail to be repeated. **LEARNED** is the line that separates an archive from a log: the exception, the quirk, the warning — the setting that looks wrong and is deliberately so, the “correction” that must not be made. It is empty on most orders and priceless when it is not, because it is where located knowledge [1] enters the permanent record. **NEXT** records watch items, follow-up work and parts on order — the line that spawns the next work order instead of relying on someone's memory to. Dictated at the asset, the five lines take about ninety seconds; typed, about three minutes. Section 9 counts those minutes without apology.

THE ARCHIVE AS IT IS	THE SAME ORDER, FIVE LINES LATER
WO-24187 · AHU-04, Bldg 4 penthouse · Status: Completed <i>"Reset unit. Done."</i>	WO-24187 · AHU-04, Bldg 4 penthouse · Trip on low-temperature safety · Status: Completed FOUND Tripped on freeze-stat during morning warm-up; third event this season. Coil ΔT normal once running — fault only present on cold starts. CAUSE Outside-air damper linkage binding at ~30% travel; actuator stalls, cold slug reaches the stat before the coil loads. ACTION Cleaned and lubricated linkage; re-stroked actuator; verified minimum OA position at 15%. Run-tested through two warm-up cycles, no trip. LEARNED Freeze-stat on this unit is set 2 °F above spec deliberately, after the 2019 coil repair — do not "correct" it without checking the glycol loop first. NEXT Linkage is worn; replacement kit listed under parts. Watch the first cold snap in November.

Figure 2. The five-line anatomy, applied. The left record proves the work happened. The right record — ninety seconds of dictation — is the difference between an archive and a log: the next responder inherits the diagnosis, the fix, the standing exception and the watch item.

5.3 The Closure–Quality Audit and the Recurrence–Recoverability Matrix

The audit is the discipline's inspection routine: a monthly sample of 25 to 40 recently closed corrective orders, scored by a rotating pair of reviewers to damp scorer drift, in an hour or two of standing effort. Sampling matters — a census is neither necessary nor survivable — and so does the deliberate coarseness of the scoring, because false precision is one of the ways exercises like this die of their own weight. Each sampled order receives two judgments, each simply high or low.

Recurrence likelihood asks: will this fault, or its family, present again — on this asset or its siblings? High-recurrence items cluster where facility work clusters: seasonal changeovers, aging equipment, fleet-common components, the assets everyone already calls trouble. **Record recoverability** asks the Re-Fix question at order level: could the competent stranger re-fix from the record alone? The two judgments place every sampled order in one of four quadrants, and each quadrant carries a standing instruction.

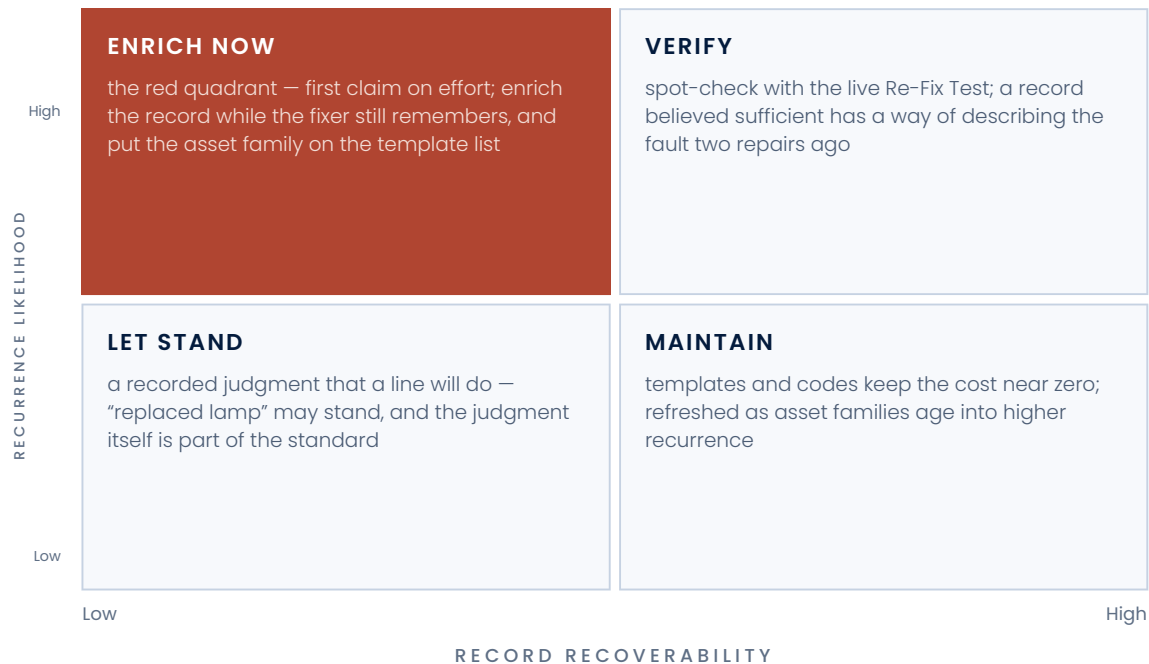


Figure 3. The Recurrence–Recoverability Matrix. Every sampled order receives one of four standing instructions; program effort concentrates in the red quadrant, where the faults most likely to return are the ones the record cannot answer.

The upper-left cell — high recurrence, low recoverability — is the audit’s reason for existing, and this paper refers to it throughout as the **red quadrant**. It is where chronic assets betray themselves: the same equipment, three closures, no content. Red-quadrant findings trigger the audit’s one intervention with a clock on it — an enrichment pass, made while the fixing technician still remembers, because a red-quadrant order enriched in the same month costs minutes and the same order reconstructed next season costs a diagnosis. Findings feed exactly three places: the template library, so the asset family’s next order opens with better prompts; the toolbox talk, where one good note is read aloud as recognition rather than correction; and the enrichment backlog, worked with the same weekly discipline as the maintenance backlog it resembles.

5.4 The capture-mechanisms menu

The discipline prescribes that the record be made; it does not prescribe a single way of making it, because technicians differ in kind as much as knowledge items do, and the right mechanism differs with them. The menu below is deliberately plural and vendor-neutral — every mechanism runs on commodity tools, and the choice among them is a matching exercise, not a procurement. Two principles govern the matching, and they are the same two this series has applied to every capture question [1]: the record must land in a system that survives its author, and the mechanism chosen must be the one the holder will actually complete. The veteran who will never type a paragraph will narrate a flawless one while pointing at the equipment; the technician allergic to prose will tap five guided prompts on a tablet at the asset without breaking stride.

Mechanism	Best suited to	Typical effort per order	Characteristic failure mode
Voice-to-text dictation at the asset	Holders who will talk but not type; multi-step fixes with physical technique	60–90 seconds	Raw transcripts filed unedited — searchable garbage without a light structuring pass

Mechanism	Best suited to	Typical effort per order	Characteristic failure mode
Photo-first closure (before / after, annotated)	Spatial fixes: linkages, labels, access, “which one, exactly” — the visual layer [1]	1–2 minutes	Images without captions; the archive knows <i>what</i> but not <i>why</i>
Guided closure form (five prompts)	Consistency across a mixed team; the anatomy as the path of least resistance	2–3 minutes; hours to configure once	Fields multiply until the form is circumvented
Problem-cause-remedy coding (small taxonomy [12])	Counting and trending: chronic assets, fault families, fleet comparisons	Seconds, once learned	“Other” absorbs the taxonomy; 400 codes where 12 would do
Template library per asset class	Routine PMs and known fault families; pre-filled prompts per equipment type	Hours per class, once	Templates pasted unedited — the note describes the template’s building, not this one
Scribe pairing	Veterans in their final years; complex one-off repairs worth a coordinator’s pen	15–30 minutes per order	Scheduled as a courtesy, cancelled under workload pressure
Escalation debrief	Emergency and after-hours work — the highest-value fixes, captured within a day	10–15 minutes	Skipped precisely on the nights that mattered

Table 1. The capture-mechanisms menu. *The discipline’s concern is that the red quadrant shrinks in systems a stranger can reach; the mechanism is matched per technician and per order, not mandated portfolio-wide.*

6. Three pillars: where the record pays back first

Closure discipline earns its keep at three kinds of moment. Each stresses the archive differently, and each converts the same five lines into a different kind of value.

6.1 The recurring fault

Diagnosis is the expensive half of corrective work, and recurrence is the norm rather than the exception: faults arrive in families — the seasonal trip, the fleet-common component, the intermittent that only presents under one weather — and the family’s second visit is where the record either pays or doesn’t. A passing note collapses the second diagnosis to a reading: the responder opens the asset history, finds FOUND and CAUSE waiting, and starts where the last fix ended instead of where the last fault began. The effect is largest exactly where operations hurt most — after hours, when the responder is whoever is on call rather than whoever knows the unit, and the note is the difference between a fifteen-minute restoration and an escalation. A working archive also ends the **reassignment chain**: the order that bounces between technicians until it finds the one person who has seen the fault before is an order searching for a record that should have existed, and its route through the CMMS is one of the clearest fingerprints of located knowledge this series has described [1]. And the coding layer converts folklore into evidence: “that unit is trouble” becomes a cause-coded trend — three linkage failures in fourteen months — which is the sentence a capital request can actually cite.

6.2 The handoff and the new hire

An archive of honest closure notes is the only operational document that grows without a project, and it is the nearest thing a building has to a syllabus. A new technician inheriting a portfolio with a working archive

onboards against the real fault history of the actual buildings: the chronic assets announce themselves, the standing exceptions are written where the equipment is, and the veteran's scarce mentoring time goes to judgment rather than to addresses. The onboarding literature suggests how much room there is to recover — the Panopto study found new hires receiving roughly 2.5 months of formal training but taking as long as six months to reach full productivity, with much of the gap spent hunting for information and reconstructing context [3] — and while a team should measure its own ramp rather than adopt anyone else's, the mechanism is not in doubt: the slowest weeks of a first year are spent rediscovering, one work order at a time, exactly what the archive should already hold. The same machinery is what makes absence survivable. Vacation coverage, the two-week resignation, the morning nobody scheduled — each is the Re-Fix Test run live, without warning, on the whole portfolio at once; and a role's Handoff Test score [1] rises with its closure archive faster than with any other single record, because the archive is where the role's daily knowledge has been accumulating one order at a time.

6.3 The contract, the warranty and the audit

The third pillar belongs to the external reader. With outsourced delivery carrying a majority of a market measured in the trillions of dollars [7], the work-order archive is, in a great many portfolios, the principal artifact one organization hands another on a changeover date — and its condition on that date is a commercial fact. Three implications follow. First, **ownership**: notes written throughout a contract term in the owner's system of record are knowledge the owner actually owns, whatever happens commercially — which is an argument for wiring the discipline into the contract rather than requesting it at the exit. Closure standards belong in service agreements as a specified deliverable — the anatomy named, the systems named — verified by exactly the acceptance sampling of Section 5.3, extending to the operating phase the structured-handover logic the construction world encoded in COBie a generation ago [13]. Second, **evidence**: warranty claims, insurance recoveries and disputed handovers are decided on contemporaneous records, and the note written the day of the failure outweighs the recollection assembled for the claim — an archive of "Done." concedes every such argument in advance. Third, **the audit**: regulators, certifiers and insurers increasingly ask not whether maintenance was scheduled but whether its execution can be evidenced, and the closure record is where execution lives. The capital-planning reader completes the set: a deferred-renewal backlog of the scale the sector reports [6] is prioritized credibly only by teams that can show, from documented failure history, which assets are consuming the operation — and the coding layer of Section 5.2 is what makes that showing a query rather than a memory exercise.

7. Implementation: crawl, walk, run

Everything in Sections 2 through 6 can be adopted incrementally, and should be. The staging below deliberately mirrors the maturity ladder of this research series [1][2]: prove the mechanism at small scale against a measured baseline, extend it under written rules, then embed it in the routines that already run the operation.




Figure 4. The implementation path. Each stage is complete in itself: a team can stop after the crawl with a measured result, or carry the same instruments to portfolio scale.

Crawl — a bounded pilot, 60 to 90 days. Scope to one trade, one team or one building cluster. Baseline first, and quietly: sample a hundred recently closed orders and score them against the Re-Fix Test *before* the standard is announced, because announcement contaminates the baseline the pilot's verdict depends on. Establish governance in its minimal viable form before the first toolbox talk — a one-page charter naming the owner, the evaluation firewall of Section 10, and what quality findings will and will not be used for, said in plain language. Then introduce the anatomy *with the crew's edits*: an hour's working session in which the technicians rename the lines into their own words and set the three-line and one-line thresholds is not a courtesy; it is the difference between a form and a discipline. Let the crew choose one capture mechanism from Table 1. Spend ten minutes of the existing weekly toolbox talk reading one good note aloud — recognition, not correction. End the pilot with a live Re-Fix Test on one genuinely closed order, and report three numbers: the baseline, the movement, and the cost in minutes. A pilot that cannot beat its baseline is treated as a finding, not a failure — it will usually mean the scoring flattered the archive, which is itself worth knowing before scaling.

“A note standard the technicians amended is a standard the technicians will use.”

Walk — portfolio extension under written cadence. Extend by asset criticality rather than alphabetically — the chillers before the ceiling tiles — at whatever pace the audit capacity can honestly sustain. Put the cadence in writing: monthly samples, rotating scorer pairs, findings routed to templates, talks and the enrichment backlog. Roll out the coding layer small: a taxonomy of twelve problem-cause-remedy codes, pruned annually, outlives a taxonomy of four hundred, and the “Other” share is its health gauge. Wire the event triggers to the places events first become visible: a chronic-asset flag routes the family into the red-quadrant enrichment sprint; a resignation or retirement notice routes the leaver's most-touched assets into a note review while questions can still be asked — the closure archive's contribution to the continuity problem this series treats in full elsewhere; a contract entering its final year starts acceptance sampling of the provider's closures against the contracted standard.

Run — embedded operations. The discipline stops being an initiative when its verbs appear inside other documents' sentences. Closure standards appear in vendor agreements as a specified, sampled deliverable and in project closeout beside the as-builts. Templates appear in the CMMS configuration as the default path, so the anatomy is the line of least resistance rather than an extra step. The Re-Fix score appears wherever the organization already reviews maintenance indicators [13][14], and the audit's hours appear in the operating

budget as routine rather than initiative funding. Onboarding curricula cite the archive. At this stage the discipline also passes its own test: its method documented, its audit owner's successor named, its history versioned — a note standard that would not survive its founder is a finding about the standard.

8. A maturity model for the work-order record

Teams need a way to locate themselves before they can plan a route, and a five-level model — parallel in spirit to the maturity ladders of this series' earlier papers [1][2] — serves the purpose. The levels describe the practice, not the people: a Level 0 organization can be staffed entirely with excellent technicians, which is in fact the most common way to be Level 0, since excellence on the payroll is what makes an empty archive painless right up until it is expensive.



Figure 5. Five maturity levels. *The levels describe the practice, not the people — a Level 0 team can be staffed entirely with excellent technicians.*

Level	Name	What it looks like	Characteristic exposure
0	Status-only	“Closed” is the record; the archive proves work happened and holds nothing else; every question routes to a person.	Every recurrence is a re-diagnosis; every departure devalues the archive to a list of dates.
1	Anecdotal	Good notes exist wherever good writers happen to work; unindexed, unowned, invisible to anyone who doesn't know to look.	Quality leaves with its authors; the archive's value is an accident of staffing.
2	Standardized	The anatomy is defined and trained; a pilot has beaten its measured baseline; capture mechanisms are matched to the crew.	Standard decays without cadence; enrichment is opportunistic.
3	Managed	Monthly audits operate; Re-Fix scores are trended; codes are clean and pruned; feedback reaches templates and talks.	Program depends on its owner; not yet written into contracts and closeout.
4	Embedded	Closure standards live in vendor contracts, project closeout, onboarding and the KPI review; the discipline survives its founder.	Residual: judgment and craft still transfer only through mentoring time.

Table 2. A five-level maturity model for the work-order record. *Most organizations assessing themselves honestly land at Level 0 or 1.*

Most facility organizations assessing themselves honestly will land at Level 0 or Level 1, and the model implies no embarrassment in that — the practice it describes has not historically existed as a discipline, so its absence is the industry's default rather than any team's negligence. The model's value is directional. The move from 0 to 2 is the pilot of Section 7 and is achievable inside a quarter. The move from 2 to 3 is where most programs stall, because it is the move from having announced a standard to operating one — cadence, sampling, feedback — and it is sustained by exactly the ownership and governance that Properties four and five of Section 4 insist on. Level 4 is not a destination so much as a disappearance: the discipline stops being visible as a program because its components have been absorbed into contracts, configurations, onboarding and the operating calendar, which is where record-keeping practices go when they succeed.

9. Measuring the program: baselines before benchmarks

This series has been consistent about measurement discipline, and the position bears restating in full: the correct benchmark is the team's own baseline, measured before the pilot begins and again after the discipline lands. This paper publishes no industry benchmark for any metric below, declines to invent one, and regards any program justified by someone else's numbers as unjustified. The published figures of Section 3 establish that the problem is real and widely shared; only the team's own instruments can establish what it costs *here* and whether the discipline is repaying its minutes. Every metric below can be captured from the CMMS or a two-week sampling exercise, and each should be baselined from day one of the crawl stage — afterwards is too late, because the standard contaminates the baseline it was supposed to be measured against.

Metric	What it measures	How to baseline it
Re-Fix score	Share of sampled closed corrective orders from which a competent stranger could re-fix, judged against the record alone; optionally criticality-weighted	Score 100 recent closed orders before the standard is announced; confirm with one observed live test
Repeat-visit rate	Corrective orders on the same asset and fault family within 90 days — the recurrence the record either answers or doesn't	CMMS query over the trailing year, coded by a coordinator
Note completeness	Share of sampled notes carrying the anatomy appropriate to the order's tier — five lines, three, or one	Same monthly sample as the audit; reported as a trend, never as a league table
Reassignment chain length	Average hand-offs an order makes before reaching resolution — the archive's failure to answer, measured in routing	CMMS assignment history over the trailing two quarters
Ask rate	Share of sampled located-knowledge queries resolved by the most-consulted people rather than by the record	Two-week self-kept tally of questions arriving by radio, phone and doorway [1]
Code hygiene	Share of failure codes landing in "Other" — the taxonomy's health, and its early warning	From the coding fields once the taxonomy operates; reviewed at the annual prune

Table 3. Program metrics. No industry benchmark is offered for any row; the benchmark is the team's own baseline, measured before the pilot and again after.

Two of the six deserve elaboration. The **ask rate** is the program's early-warning gauge, and the metric most likely to be misread as being about the people consulted rather than the systems that fail to answer — which is why it is reported at team level only, and why Section 10's firewall exists. A falling ask rate is the discipline working: questions migrating from doorways to records. The **code-hygiene** share is the canary of taxonomy death: when "Other" passes a quarter of entries, the taxonomy has stopped describing the operation and started defending itself, and the annual prune is due early. And a caution that applies to every row: closure percentages computed against a stale sample measure a team that no longer exists, so the audit's recency — the share of the trailing quarter actually sampled — is reported beside every other number, the way a laboratory reports the calibration date next to the reading.

An illustrative cost model. To show the shape of the economics — and only the shape — consider a deliberately assumed case, every input of which is a stated assumption rather than a benchmark, a promise

or a finding: a 20-technician team across a 25-building portfolio, closing roughly 6,000 work orders a year, running the crawl and walk stages in one year with two hires expected.

An illustrative year-one model — every input a stated assumption	Hours
<i>Assumed case: 20 technicians; 25 buildings; ~6,000 closable work orders per year; two hires expected; crawl and walk stages completed in year one.</i>	
Effort — design, charter, anatomy workshop and template setup	20
Effort — rollout training and toolbox-talk facilitation (20 technicians at ~2 h combined)	40
Effort — capture-mechanism setup (dictation workflow, photo path, CMMS field configuration)	24
Effort — monthly Closure-Quality Audits (2.5 h × 12, rotating pairs)	30
Effort — the three-minute tax: 2.5–3 minutes of capture on ~6,000 closures	250–300
Effort — red-quadrant enrichment passes (~40 orders at ~1 h, while the fix is fresh)	40
Year-one effort, total	~405–455
Recovered — diagnosis time on recurring-family faults (a fifth of orders recur in family; 30 minutes saved on a third of them)	~200
Recovered — interrupt load redirected to the record (4 located-knowledge queries per working day at ~0.5 h combined asker-and-answerer cost; one quarter redirected)	~125
Recovered — repeat visits avoided (1 per building per quarter at 1.5 h; half eliminated)	~75
Recovered — new-hire ramp reduction (2 hires; 2 weeks shorter at a 50% effectiveness gap)	~80
Recovered — one changeover, warranty claim or dispute carried by the record rather than by recollection (conservative)	~120
Year-one recovery, before any avoided escalation, failed audit or lost claim	~600

Table 4. An illustrative cost model. *The inputs are assumptions, not benchmarks or promises; the point is the shape — bounded, predictable effort against open-ended, compounding exposure. Teams should rebuild the table with their own baselines.*




Figure 6. *The shape of the illustrative model: bounded, predictable effort against open-ended, compounding exposure. Inputs are assumptions — teams should redraw both bars from their own baselines.*

“The discipline’s largest cost arrives three minutes at a time. Its critics will count those minutes — and the model counts them too.”

On these assumptions, the first year spends roughly 405 to 455 hours and recovers on the order of 600 — before counting a single avoided emergency escalation, failed audit finding, conceded warranty claim or changeover absorbed unprepared, which is where the sums that dwarf this table actually live. The model's honesty matters more than its outputs: the largest line on the effort side is the capture tax itself, permanent and unhideable, and the discipline's case must be made with that line on the table. What the shape shows is that the effort is bounded and predictable — it scales with order volume and shrinks after year one to cadence maintenance — while the costs it offsets are open-ended and compounding, growing with every recurrence, departure and transition the calendar delivers. A team should rebuild the table with its own inputs, and a team whose rebuilt table cannot beat its baseline has learned something real: either its archive is better than it feared, or its scoring flattered it — and one live Re-Fix Test will say which.

10. Governance: a record about assets, not a dossier on people

A closure archive gathers information from and about individuals: authorship, timestamps, and — when the discipline is working — honest accounts of what was found, including what a previous repair got wrong. Handled carelessly, the same exercise that protects the organization corrodes the honesty it depends on. A technician who suspects the notes are a productivity instrument will write notes fit for a productivity instrument: short, safe and content-free. Governance is therefore not a compliance appendix to the discipline; it is a condition of getting true records, and this paper treats it as non-negotiable in the same spirit as its predecessors treated the governance of building imagery and knowledge inventories [1].

“The day the closure note becomes a stopwatch is the last day it tells the truth.”

Five rules constitute the minimum. **First, the firewall:** note content, note quality and note timestamps feed no performance evaluation, no compensation decision and no pace analysis — stated in the charter, said aloud in the rollout, and kept, because the first exception ends the program's supply of honesty permanently. The audit measures the state of the organization's record, and a sample full of empty notes is a finding about the standard, the templates and the mechanisms — never about the people who closed what the system made it easy to close emptily. Quality findings are reported at team level; recognition is individual, correction is structural. **Second, just-culture handling of error narratives.** The most valuable notes in any archive are the ones that record what went wrong — the bypassed interlock found in place, the prior fix that caused the fault, the near-miss. Those sentences get written only where safety practice's hard-won rule applies: the report of an error is not the confession of one, and the organization's response is to fix the condition, not to find the author. An archive that punishes candor will be sanitized into uselessness within a quarter, and no one will be able to say when it happened. **Third, attribution with dignity.** Author fields exist for reachability and for credit — the follow-up question during a transition, the good note read aloud at the toolbox talk — not for league tables of who writes most or is asked most; the consultation patterns the ask rate samples are reported in aggregate for exactly this reason. **Fourth, data-protection by design.** Notes and their photographs contain names, tenant information and images of occupied spaces; they are personal data within the meaning of regimes such as the EU's General Data Protection Regulation [15], and they inherit the imagery rules this series has already set out [1]: minimization to operational necessity, stated purpose, rule-based retention, role-based access, and revocation wired into offboarding. **Fifth, honest workload accounting.** The minutes are work: planned into job standards, visible in wrench-time expectations, and credited — with scribe support provided rather than implied for the technicians who need it. A discipline funded by pretending its cost is zero will be paid for out of honesty instead.

None of these rules is costly. All of them are cheaper than an archive of fiction.

11. Standards and frameworks alignment

Closure discipline is not a standard and does not compete with one; it is a practice that gives several existing frameworks the operational instrument they assume but do not specify. The alignments below are offered so that teams operating under these frameworks can position the work inside obligations they already carry.

Framework	Where the discipline fits
ISO 41001:2018 — Facility management systems [16]	Operationalizes the documented-information and operational-planning clauses at the work-order layer; audit samples and Re-Fix scores serve as objective evidence of control.
ISO 55000 series — Asset management [17]	Supplies the operating-phase asset information the standards require: condition, failure history and intervention records attached to the assets they describe, feeding asset-risk registers and renewal cases.
ISO 30401:2018 — Knowledge management systems [18]	A facility-native mechanism for capturing operational knowledge as a managed resource — at the point of creation, under cadence and ownership.
ISO 14224:2016 — Reliability and maintenance data [12]	Source of the problem-cause-remedy taxonomy logic, imported from the reliability world and deliberately cut down to facility scale: codes for counting, prose for knowing.
EN 15341:2019 — Maintenance KPIs [13]	Supplies the indicator discipline the Section 9 metrics slot into; the Re-Fix score and repeat-visit rate sit naturally beside established maintenance indicators.
SMRP best-practice metrics [14]	Aligns the program's measures with the maintenance-and-reliability profession's metric conventions and definitions.
IFMA competency framework [19]	Maps to operations and maintenance; facility information management and technology management; performance and quality; risk management; and communication.
COBie / BIM handover conventions [20]	Extends structured-handover logic from commissioning data into the operating phase: contracted closure standards mirror COBie's deliverable discipline at every service transition.
GDPR and data-protection regimes [15]	Governs the people-adjacent content of notes and photographs: minimization, stated purpose, rule-based retention and access control.

Table 5. Standards and frameworks alignment.

The alignment with ISO 41001 and the ISO 55000 series deserves the emphasis. Both frameworks require documented information and asset information adequate to the organization's operations, and both are routinely experienced, at the plant-room layer, as abstractions [16][17]. The closure archive is the concrete, facility-native way those requirements are met where they matter most: the inventory of what the assets have actually done, gathered continuously by the only process that touches every asset on a schedule — the work order itself. An organization pursuing or holding certification will find the discipline slots into its support and planning clauses as the missing operational layer between the requirement to maintain documented information and the plant-room reality that requirement is ultimately about; ISO 14224's contribution, meanwhile, is deliberately partial — the taxonomy logic travels, the petroleum-scale taxonomy does not, and the discipline's twelve-code counsel is the difference [12].

12. Limitations and outlook

Limitations. This paper's claims are bounded, and the boundaries deserve the same prominence as the argument. The note captures the locatable layer of a fix, not craft. Diagnostic intuition — the feel for a bearing's sound the week before it fails — is developed, not documented; philosophers of knowledge have long observed that people know more than they can tell [11], and the discipline's contribution to that layer is indirect: it frees mentoring time from address-transfer so the time can go to judgment-transfer, and any program promising to bottle the veteran is overpromising by roughly the amount that matters most. The archive is evidence, not truth: self-reported records can be wrong, and the live test and the sampled audit exist to catch drift, not to abolish it. Scoring is structured judgment, not measurement — two reviewers can judge the same note differently, the defenses are coarse scales, rotating pairs and the live test as arbiter, and the residual subjectivity is a reason to track movement rather than worship absolute values. The three-minute tax is real and permanent, and a team that cannot fund minutes should not pretend to fund a program. Taxonomies ossify — the annual prune is load-bearing, and the "Other" share is its gauge. And the discipline scales with criticality, not with everything: a small team with genuine redundancy may honestly hold most of its portfolio at a one-line standard and be right — the judgment, recorded, is itself the discipline operating at exactly the scale the risk deserves.

Outlook. Three developments will change the discipline's mechanics without changing its logic. Speech-to-text and language models are already collapsing the cost of the capture step — turning a ninety-second narration at the asset into a structured five-line draft for the technician to correct and sign, with the signature staying human, because a record no one vouched for is not evidence. Text mining over CMMS history will surface the red quadrant as a standing dashboard rather than a monthly sample: chronic families, empty-note hotspots and reassignment chains flagged as they form. And the convergence this series has traced elsewhere [2] will pull the note toward the building's spatial record, so that the fix, its photograph and the place they describe resolve to the same address. Every one of these developments makes capture cheaper, and none of them supplies an anatomy, a cadence, a firewall or the honesty the firewall protects — which is to say the durable investment remains what it has been throughout this series: not any particular tool, but the practice.

13. Conclusion

Facility organizations have always known, in the way that everyone knows and no one budgets for, that their buildings run on knowledge that lives in people. What has been missing at the work-order layer is not awareness but instrumentation: an anatomy that makes the record definable, a test that makes it measurable, an audit that makes it improvable, mechanisms that make it completable by the people who hold the knowledge, and the governance that makes the whole exercise answerable with the truth. The closure note is that instrumentation, applied at the one moment ordinary operations already deliver — every day, at every asset — the problem, the solution and the person who understands both.

The practice asks little that facility teams do not already know how to do. It is, in structure, preventive maintenance applied to the organization's memory: an inspection on a cadence, findings in a backlog, standards with owners, verification by test — performed on an asset class the CMMS has no module for and the balance sheet has no line for, which nonetheless determines what every fault costs the second time. The work order already exists. The technician already stands at the asset. The knowledge is already in hand, complete and fresh, for the length of one open text field. The building will present the fault again — the season will turn, the linkage will bind — and the only genuinely open question is what the next responder will find waiting: five lines, or one word.

References

- [1] Hanov, N. *The Visual Asset Register: How 360-Degree Documentation Improves Work Orders, Inspections and Staff Handoffs*. Levaru, 2026. levaru.co/white-papers/visual-asset-register/
- [2] Hanov, N. *Minimum Viable Digital Twin for Existing Buildings: A Facility Manager's Practical Framework for 360° Reality Capture, Structured Asset Data and Operational Integration*. Levaru, 2026. levaru.co/white-papers/minimum-viable-digital-twin/
- [3] Panopto and YouGov. *Workplace Knowledge and Productivity Report* ("Valuing Workplace Knowledge"). Panopto, 2018. Survey of more than 1,000 United States workers; figures cited: 42 percent of institutional knowledge unique to the individual; 5.3 hours per employee per week waiting for or recreating information; approximately 2.5 months of formal new-hire training against up to six months to full productivity.
- [4] McKinsey Global Institute. *The Social Economy: Unlocking Value and Productivity Through Social Technologies*. McKinsey & Company, 2012. Figure cited: interaction workers spend approximately 19 percent of the workweek searching for and gathering information.
- [5] Quickbase. 2023 survey on workplace information friction; figure cited: 53 percent of respondents spend more than 10 hours per week chasing information. As discussed in [1].
- [6] Gordian. *The State of Facilities in Higher Education*. Gordian (annual). Figure cited: deferred capital renewal on North American campuses of US\$156 per gross square foot.
- [7] Grand View Research. *Facility Management Services Market Size, Share and Trends Analysis Report*. Grand View Research, 2025. Figures cited: global market of approximately US\$1.75 trillion in 2024; outsourced segment accounting for 61.5 percent of global revenue in 2024.
- [8] Study cited in the International Facility Management Association's *Facility Management Journal* finding that 31 percent of facility management employees are older than 55 and eligible to retire between 2022 and 2037; discussed in [1].
- [9] University of Tennessee, Knoxville, Facilities Services: approximately 600 staff maintaining more than 250 buildings, with 97 retirements recorded between 2020 and 2025; as discussed in [1].
- [10] DeLong, D. W. *Lost Knowledge: Confronting the Threat of an Aging Workforce*. Oxford University Press, 2004.
- [11] Polanyi, M. *The Tacit Dimension*. University of Chicago Press, 1966.
- [12] ISO 14224:2016. *Petroleum, petrochemical and natural gas industries — Collection and exchange of reliability and maintenance data for equipment*. International Organization for Standardization, Geneva. Cited for its problem-cause-remedy taxonomy logic, adapted here to facility scale.
- [13] EN 15341:2019. *Maintenance — Maintenance Key Performance Indicators*. European Committee for Standardization, Brussels.
- [14] Society for Maintenance & Reliability Professionals. *SMRP Best Practices*, metrics and guidelines for maintenance and reliability indicators. SMRP, Atlanta.
- [15] Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation).
- [16] ISO 41001:2018. *Facility management — Management systems — Requirements with guidance for use*. International Organization for Standardization, Geneva.
- [17] ISO 55000 series. *Asset management*. International Organization for Standardization, Geneva.
- [18] ISO 30401:2018. *Knowledge management systems — Requirements*. International Organization for Standardization, Geneva.
- [19] International Facility Management Association. IFMA core competencies of facility management (competency framework).
- [20] COBie — Construction-Operations Building information exchange, as maintained within the United States National BIM Standard (NBIMS-US), buildingSMART alliance.



LEVARU

Research on the operational record of buildings

Also in this series

The Visual Asset Register · 2026

Minimum Viable Digital Twin for Existing Buildings · 2026

The Continuity Audit · 2026

levaru.co/white-papers · info@levaru.co · +1 (703) 646-8300