




LEVARU WHITE PAPER

The Continuity Audit

A Facility Manager's Framework for Measuring and Closing Located-Knowledge Risk Before Retirements, Contract Changes and Reorganizations

Nikolai Hanov

Published by Levaru · August 2026



LEVARU

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

RESEARCH SERIES · PAPER 3

Abstract

Facility organizations run on **located knowledge** — knowledge tied to specific places, assets and relationships that is real, operational and almost never written down. That knowledge leaves through three foreseeable doors: retirements remove its holders outright, contract changes remove them wholesale on a changeover date, and reorganizations remove them quietly while they remain on the payroll. The standard responses — the exit interview, the farewell-quarter binder — arrive too late and capture too little. This paper defines the **Continuity Audit**: a governed, repeatable assessment that inventories the located knowledge an operation depends on, scores each item for operational criticality and documentation coverage, and converts the result into an owned, prioritized closure plan — re-run on a defined cadence and re-triggered by the events that put knowledge at risk. It supplies the instruments — an expanded Handoff Test, a located-knowledge inventory method, the Criticality–Coverage Matrix and a vendor-neutral menu of closure mechanisms — together with a staged implementation path, a five-level maturity model, a baseline-first metrics discipline, an explicitly illustrative cost model, and the governance that keeps an audit of knowledge from ever becoming an audit of people. It aligns the practice with ISO 30401, ISO 41001, the ISO 55000 series, COBie handover conventions and IFMA’s competency framework, and closes with candid limitations.

KEYWORDS

knowledge continuity

located knowledge

facility management

workforce transitions

succession

outsourcing transitions

onboarding

ISO 30401

ISO 41001

Suggested citation: Hanov, N. (2026). *The Continuity Audit: A Facility Manager’s Framework for Measuring and Closing Located-Knowledge Risk Before Retirements, Contract Changes and Reorganizations*. Levaru. levaru.co/white-papers/continuity-audit/

About this series: This is the third paper in Levaru’s research series on the operational record of buildings, following *The Visual Asset Register* (January 2026) and *Minimum Viable Digital Twin for Existing Buildings* (2026).

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

Contents

1. The problem: knowledge with an address	4
2. Definition: what a Continuity Audit is — and is not	5
3. Why now: three converging pressures	7
4. What makes it an audit, not a survey: five properties	9
5. The instruments	10
5.1 The Handoff Test	10
5.2 The Located-Knowledge Inventory	11
5.3 The Criticality–Coverage Matrix	12
5.4 The closure-mechanisms menu	13
6. Three pillars: where the audit pays back first	14
6.1 Retirements and planned departures	14
6.2 Contract changes and vendor transitions	15
6.3 Reorganizations, role changes and unplanned absence	16
7. Implementation: crawl, walk, run	16
8. A maturity model for knowledge continuity	18
9. Measuring the program: baselines before benchmarks	19
10. Governance: auditing knowledge, not people	21
11. Standards and frameworks alignment	22
12. Limitations and outlook	23
13. Conclusion	24
References	25

1. The problem: knowledge with an address

Every facility team runs on a category of knowledge that has no good name in most organizations, though everyone recognizes it the moment it is described. It is the knowledge of which breaker panel feeds the east wing and which one merely looks like it does; of the valve that isolates the third-floor riser and the trick to closing it without hammering the line; of why the air handler on the roof of Building 4 was piped backwards in 1998 and what happens to the loop if anyone "corrects" it; of which tenant contact actually answers the phone at 2 a.m. and which door the fire marshal always checks first. Levaru's earlier research called this **located knowledge**: knowledge that is real, operational and tied to specific places, assets and relationships — and almost never written down, because prose is a poor medium for it and because, for as long as its holder is one radio call away, it never needs to be [2].

“For as long as its holder is one radio call away, located knowledge never needs to be written down — until, on a known date, it does.”

Located knowledge is not the same as skill. A licensed electrician carries their trade from building to building; what they cannot carry in is thirty years of a particular building's decisions, exceptions and workarounds. Philosophers of knowledge have long distinguished what people can articulate from what they simply know how to do — Polanyi's observation that human beings know more than they can tell [13] — but the facility version of the problem is more specific and more tractable than tacit skill. Most located knowledge is not ineffable. It is perfectly expressible — someone standing in the room can point to the panel and say "that one, not the one with the newer label" — and it goes unrecorded not because it cannot be captured but because no one has ever been assigned to capture it, no system has a field for it, and the organization has never measured how much of it there is or what it would cost to lose.

The loss, when it comes, arrives through three doors. A **retirement or resignation** removes a holder outright, and with them every answer they were never asked to write down. A **contract change** — an outsourcing transition, an insourcing decision, a re-tender that replaces one service provider with another — can remove a dozen holders on a single changeover date, none of whom worked for the building owner and some of whom are contractually discouraged from lingering. A **reorganization** removes holders more quietly: the technician is still employed, but now covers a different portfolio, reports through a different structure, or has been promoted away from the tools, and the fifteen-minute answer becomes a favor that decays into a bottleneck and then into silence.




Figure 1. Three foreseeable doors. Retirements, contract changes and reorganizations differ in mechanics but share one property: the date is known in advance, and located knowledge leaves through all three.

What these three events share is that they are, in most organizations, entirely foreseeable and almost entirely unprepared for. Retirement dates are known months or years in advance. Contract end dates are written into the contracts themselves. Reorganizations are planned in conference rooms long before they are announced. Yet the standard organizational responses — the exit interview, the transition binder assembled in the final fortnight, the hopeful instruction to "document your processes before you go" — consistently arrive too late, capture too little, and capture the wrong things, because they are improvised at the moment of departure rather than operated as a practice in the years before it.

This paper proposes that practice. It defines the **Continuity Audit**: a governed, repeatable assessment that makes located-knowledge risk visible, comparable and closable before the events that realize it — and it argues that for most facility organizations, the audit is not a human-resources nicety but an operational control, as fundamental as the preventive-maintenance schedule and considerably cheaper than the failures it prevents.

2. Definition: what a Continuity Audit is — and is not

Definition. A **Continuity Audit** is a governed, repeatable assessment that inventories the located knowledge a facility organization depends on, scores each item for operational criticality and for documentation coverage, and converts the result into an owned, prioritized closure plan — re-run on a defined cadence and re-triggered by departures, contract changes and reorganizations, so that the risk picture stays as current as the team it describes.

Each clause of the definition is load-bearing. *Governed* means the audit has a named owner, documented rules for how personal information is handled, and access controls on its outputs — because a candid map of where an organization's knowledge is concentrated is also a map of its fragility, and it deserves the same protection as any other sensitive operational record. *Repeatable* means the method is written down and survives its facilitator, so that the third audit is comparable with the first. *Inventories* means the unit of analysis is the knowledge item, not the person: "isolation sequence for the Building 2 chilled-water loop," not "everything Maria knows." *Scores* means every item receives two judgments — how much it matters and how findable it is without its holder — so

that effort can be ranked rather than merely listed. *Closure plan* means the audit ends in assigned work with dates, not in a report. And *cadence* means the audit is an operating practice, not a project: an inventory that is not refreshed decays at exactly the rate the team changes, which is the rate that made the audit necessary in the first place.

The Continuity Audit sits deliberately between two established responses, and it is defined as much by what it declines from each as by what it takes. At one pole is the **exit interview** and its cousins — the handover memo, the last-month transition binder. These are cheap, familiar and almost always too late: they compress years of located knowledge into the busiest weeks of a departure, they depend entirely on what the departing person happens to recall under prompting, and they capture answers to questions no one yet knows to ask. At the other pole is the **enterprise knowledge-management program** — the platform procurement, the taxonomy project, the mandate that every team maintain its wiki. These are systematic but heavy, they are designed around office knowledge work rather than plant rooms, and in facility organizations they routinely stall at exactly the layer that matters most, because the knowledge that lives behind the locked mechanical-room door was never going to arrive by asking technicians to write articles. The Continuity Audit is far more rigorous than the former and far lighter than the latter: it borrows the audit disciplines of scope, scoring, cadence and evidence, and applies them to a domain the enterprise programs rarely reach.

Two boundary clarifications complete the definition. First, the audit assesses *knowledge*, not *people*. It produces no rating of any individual's performance, value or replaceability; it asks only where operationally necessary knowledge resides and whether the organization could function without asking. Section 10 treats this boundary as a governance obligation rather than a stylistic preference. Second, the audit is a measurement and planning instrument, not itself a documentation method. What closes a red-quadrant item might be an annotated 360-degree capture, a one-page runbook, a narrated walkdown video, a structured shadowing rotation or a corrected CMMS record; the audit's job is to determine *which items deserve that effort, in what order, and whether the effort worked* — a question that, in the absence of an audit, is answered by whoever retires next.




Figure 2. The Continuity Audit as an operating loop. Six steps run on a defined cadence; any of the three trigger events restarts the loop early, while preparation is still possible.

3. Why now: three converging pressures

None of the pressures behind the Continuity Audit is new. Buildings have always outlived careers, and every generation of facility leadership has told a version of the story about the technician who retired with the building in their head. What has changed is the arithmetic: the rate of holder departure, the price of information friction and the share of facility work delivered through contracts that periodically change hands have all moved in the same direction at once. Three published figures — each cited from independent sources rather than generated by this paper — frame the moment.

The first pressure is workforce exit at scale. 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 [1] — nearly a third of the profession's holders of located knowledge, on a countdown that is already several years advanced. The pattern is visible wherever institutions publish their numbers: at the University of Tennessee, Knoxville, a facilities organization of roughly 600 staff maintaining more than 250 buildings recorded 97 retirements between 2020 and 2025 [3] — an average of one departing veteran roughly every three weeks, sustained for half a decade, in a single portfolio. Research on aging workforces has warned since at least DeLong's *Lost Knowledge* in 2004 that organizations systematically underestimate what leaves with long-tenured staff [12]; two decades later, the facility profession is living inside that warning's timeline.

The second pressure is the measured cost of knowledge that lives in people. 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 — available to colleagues only by asking, and to no one at all once the holder is gone [4]. The same study 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 [4]. These figures describe knowledge work generally; facility work adds a physical dimension, because the waiting happens on ladders and the recreating happens through repeat site visits. They are consistent with the broader information-friction literature this research series has cited before: McKinsey's estimate that interaction workers spend roughly 19 percent of the workweek searching for and gathering information [5], and a 2023 Quickbase survey in which 53 percent of respondents reported chasing information more than 10 hours a week [6]. A facility team does not need to accept any of these numbers as its own — Section 9 insists on measuring the team's actual baseline — but their consistency across independent studies establishes the shape of the problem: unwritten knowledge is not free storage; it is a standing tax, paid in interruptions, and it converts into permanent loss on the holder's last day.

The third pressure is that facility work increasingly changes hands by contract. 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 [8]. Whatever the precision of any single market estimate, the structural fact is not in dispute: in a majority of the market, the people who hold a building's located knowledge are employed not by its owner but by a service provider — and service contracts end. Every re-tender, every insourcing decision, every provider consolidation is a scheduled knowledge-loss event with a date on it, and the knowledge at risk frequently belongs, in the practical sense, to no one: the outgoing provider has no incentive to invest in its transfer, the incoming provider does not yet know what to ask for, and the owner discovers the gap in the first winter after changeover. Deferred-maintenance pressure sharpens the stakes — Gordian's *State of Facilities in Higher Education* puts deferred capital renewal on North American campuses at US\$156 per gross square foot [7] — because a backlog of that size is navigable only by teams that know their buildings' actual condition, history and quirks, which is precisely the knowledge these transitions put at risk.

THE PRESSURE, IN PUBLISHED FIGURES

31%

of facility management employees are older than 55 and eligible to retire between 2022 and 2037 [1]

97

retirements in five years in one 250-building university portfolio — one veteran every three weeks [3]

5.3 h

lost per employee per week waiting for, or recreating, knowledge a colleague holds [4]

61.5%

of a US\$1.75 trillion global facility-services market is outsourced delivery — and contracts end [8]

Sources: IFMA Facility Management Journal [1]; University of Tennessee, Knoxville [3]; Panopto / YouGov [4]; Grand View Research [8]. Full citations in the References.

Three doors, one draft. Retirements remove holders permanently, contract changes remove them wholesale, reorganizations remove them quietly — and the organizations most exposed are the ones that cannot say, with evidence, how exposed they are. Producing that evidence is the audit's first job.

4. What makes it an audit, not a survey: five properties

Many organizations have gestured at this problem before. There have been skills matrices that recorded certifications but not building knowledge; "brain-dump" sessions scheduled in a retiree's final month; succession plans that named a successor without transferring what the successor would need; and staff surveys that asked people what they knew and filed the answers. The Continuity Audit is distinguished from these gestures by five properties. Remove any one of them and the exercise reverts to a survey — interesting once, stale within a year, and absent when the departure actually lands.

Property one: a scored inventory, not a list. The audit's core artifact is an inventory of located-knowledge items, each scored on two independent axes — operational criticality (what happens, and how fast, if this knowledge cannot be reached when needed) and documentation coverage (how much of it a competent stranger could recover from the organization's systems without asking anyone). A list of knowledge items flatters everyone and prioritizes nothing. Scoring forces the uncomfortable, useful judgments: that the isolation sequence for the data-center cooling loop and the location of the holiday decorations are both "things Maria knows," and that only one of them belongs in the first closure sprint. Section 5.3 gives the scoring instrument.

Property two: named ownership and sign-off. The audit has an owner — a named role, not a committee — accountable for its cadence, its governance and the working of its closure backlog. Each closure item likewise carries a named owner and a date. Facility teams already understand why this matters: an inspection program without an owner is a clipboard in a drawer, and a knowledge inventory without one is a spreadsheet with the same destiny. Ownership is also what makes the audit survive its founding enthusiast, which is itself a continuity question the audit should be able to pass.

Property three: coverage mapped against systems of record. For every item, the audit records not only whether documentation exists but where a stranger would find it — in the CMMS or IWMS record, in the drawing set, in a standard operating procedure, in a visual asset register [2], in a vendor manual, or nowhere. Mapping coverage to named systems does two things a yes/no answer cannot. It exposes the difference between documented and findable — knowledge captured in a retired employee's personal folder is, for continuity purposes, undocumented — and it generates the audit's most quietly valuable by-product: a gap map of the organization's information architecture, showing which systems the team's real knowledge actually lives in and which are ceremonial.

Property four: defined cadence and event triggers. The inventory is refreshed on a written schedule — annually for most portfolios, more often for teams in flux — and re-run in targeted form when specific events fire: a resignation notice, a retirement declaration, a contract entering its final year, a reorganization entering planning. Cadence is what separates a practice from a project. A team changes at some rate every year; an inventory refreshed at less than that rate describes a team that no longer exists. The triggers matter as much as the schedule, because the events that most need a current inventory — the ones in this paper's subtitle — are exactly the events that make one obsolete.

Property five: governed handling of people-adjacent data. Although the audit's unit of analysis is the knowledge item, its raw material is gathered from and about individuals, and its outputs reveal organizational dependencies that could be misread or misused. A real audit therefore operates under written rules established

before the first interview: what is collected, who can see which outputs, how long records are retained, how the exercise is explained to participants, and an explicit commitment — kept — that nothing gathered feeds performance evaluation. Section 10 develops these rules; the property is listed here because governance is a defining characteristic of the audit, not a compliance layer added afterwards.

What a Continuity Audit is not. The definition is completed by its exclusions, stated plainly. The audit is *not a performance review*, and the moment it is used as one it will never again receive an honest answer. It is *not an exit interview*, though it makes exit interviews dramatically better by ensuring the important questions are known years in advance. It is *not a platform purchase*: no software category is required to run one, and no software category substitutes for one. It is *not a succession plan*, a retention strategy or a compensation instrument — it will surface facts that make all three urgent, but deciding who advances, who is persuaded to stay and what that is worth are leadership decisions the audit informs and cannot make. And it is *not a promise that nothing will be lost*. Some of what a thirty-year veteran carries is judgment, and judgment does not transfer by documentation. The audit's claim is narrower and defensible: the *locatable* layer — the addresses, sequences, exceptions and relationships that make up located knowledge — can be inventoried, prioritized and substantially closed, and the organization that has done so meets every departure with a plan instead of a scramble.

A Continuity Audit is	It is not
A scored inventory of located-knowledge items — criticality × coverage, per item	A performance review of any individual, ever
Owned — a named audit owner; every closure item carries an owner and a date	An exit interview , however well it is run
Mapped to systems of record a competent stranger can actually reach	A platform purchase — no software category is required, and none substitutes
Operated on a cadence , re-triggered by departures, contracts and reorganizations	A succession plan , retention strategy or compensation instrument
Governed — evaluation firewall, access control, data-protection by design	A promise that nothing will be lost — judgment transfers through mentoring, not files

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

5. The instruments

Four instruments turn the definition into practice. None requires special software; all four require honesty.

5.1 The Handoff Test

Levaru's earlier work on visual documentation proposed a heuristic it called the Handoff Test [2]; this paper promotes it from heuristic to instrument and makes it the audit's headline measure.

The Handoff Test. If the most experienced person in a given role were unreachable for 90 days — no calls, no texts, no "just one question" — what percentage of the located knowledge that role depends on could a competent stranger recover from the organization's systems alone? The honest answer, expressed as a percentage, is the role's Handoff Test score. The distance between that score and 100 percent is the continuity risk the audit exists to close.

Applied as an instrument rather than a thought experiment, the test acquires a method. For each key role, the audit takes the role's inventory items (Section 5.2), asks for each one "could the stranger recover this from a named system?", and computes the score as the coverage-weighted share of items recoverable — optionally weighted again by criticality, so that a role whose documented items are the trivial ones scores as badly as it should. The test can also be run empirically, and occasionally should be: give a competent outsider — a technician from another site, a new hire past orientation, a trusted contractor — a routine task in an unfamiliar building, forbid them from asking people, and watch where the systems fail them. Thirty minutes of observed struggle is worth a hundred self-assessments, and teams that run the live version rarely argue with their score again.

Three properties make the test the right headline measure. It is *role-anchored*, so it survives individual departures and can be tracked across them. It is *directional*, moving in the right direction only when closure work actually lands in findable systems. And it is *bracing*: most teams running it for the first time produce scores they decline to publish, which is precisely the information the organization did not previously have.

5.2 The Located-Knowledge Inventory

The inventory is the audit's evidentiary base: a structured list of the located-knowledge items an operation depends on, gathered per role or per building. The failure mode to design against is the blank page — asking a veteran "what do you know that isn't written down?" produces either modesty or paralysis, because located knowledge is invisible to its holder in exactly the way water is invisible to fish. Elicitation therefore works from prompts rather than open questions, and four sources supply them.

The walk-through interview is the primary method: sixty to ninety minutes, on foot, through the interviewee's actual spaces, with the facilitator asking location-anchored questions — what should the next person know about this room; what here breaks in a way the manual doesn't describe; what would you check first if this system failed tonight; who do you call about this, and who calls you? Walking matters because located knowledge is indexed spatially in memory; people recall standing in front of the panel what they cannot recall at a conference table, and the pointing that accompanies the answer is itself content worth capturing. **Work-order mining** supplies the second source: the CMMS text fields, scanned for the fingerprints of located knowledge — orders resolved with no useful closure notes, orders reassigned repeatedly until they found the one person who could do them, recurring faults whose fix is never written the same way twice. **The call pattern** is the third: a simple two-week tally, kept by the most-consulted staff themselves, of the questions that reach them by radio, phone and doorway — each tally mark a located-knowledge item that exists nowhere else, pre-sorted by demand. **The incident record** is the fourth: after-action notes and near-miss reports read not for what failed but for what the response depended on knowing, since emergencies are the most expensive possible way to discover an inventory

item.

Each item enters the inventory in a consistent, deliberately small schema: what the knowledge is, where it applies (building, system, space), who currently holds it, where — if anywhere — it is documented, and the two scores of Section 5.3. A portfolio-scale inventory of a few hundred items fits in a spreadsheet, and should: the instrument's power is in the discipline of the schema, not the sophistication of the container.

5.3 The Criticality–Coverage Matrix

Scoring converts the inventory from a catalog into a decision. Each item receives two independent judgments, each on a deliberately coarse scale — high or low is enough to act on, and false precision in scoring is one of the ways exercises like this die of their own weight.

Operational criticality asks: if this knowledge could not be reached when needed, what happens, and how fast? High-criticality items sit on the paths of safety, regulatory exposure, business continuity and major asset protection — isolation sequences, shutdown procedures, life-safety system quirks, single-points-of-failure and the undocumented exceptions to all of them. Low-criticality items cost minutes and mild annoyance. **Documentation coverage** asks the Handoff Test question at item level: how much of this could the competent stranger recover from named systems alone? High coverage means a findable, current, sufficient record exists in a system that survives its author; low coverage means the item lives in memory, in a personal notebook, or in a file whose existence is itself located knowledge.

The two judgments place every item in one of four quadrants, and each quadrant carries a standing instruction:




Figure 3. The Criticality–Coverage Matrix. Every inventory item receives one of four standing instructions; program effort concentrates in the red quadrant.

The upper-left cell — high criticality, low coverage — is the audit's reason for existing, and this paper refers to it throughout as the **red quadrant**. Red-quadrant items get first claim on closure effort, named owners and dates. The upper-right quadrant gets *verification*, not congratulation: documentation believed current has a way of describing the building as it stood two renovations ago, and a spot-check — ideally the live Handoff Test — is cheap insurance against discovering the gap in an emergency. The lower-left quadrant is *accepted risk*, and the acceptance is recorded: writing down "we know this is undocumented and have decided it can stay that way" is itself a continuity artifact, because it tells the successor organization the gap was a choice rather than an oversight. The lower-right quadrant is simply *maintained* through normal operations.

5.4 The closure-mechanisms menu

The audit prescribes closure; it does not prescribe a single closure method, because located-knowledge items differ in kind and the right record differs with them. The menu below is deliberately plural and vendor-neutral — every mechanism on it can be executed with commodity tools, and the choice among them is a matching exercise, not a procurement.

Mechanism	Best suited to	Typical effort per item	Characteristic failure mode
Annotated 360-degree capture (a visual asset register [2])	Spatial items: locations, access routes, panel and valve identification — "which one, exactly"	Minutes per annotation once spaces are captured	Register lapses its re-capture cadence and drifts from the building
Short written runbook or SOP	Sequences: isolation order, startup and shutdown, seasonal changeover	1–3 hours	Written once, filed outside the system of record, never revised
Narrated walkdown video	Holders who will talk but not write; multi-step tasks with physical technique	Under 1 hour, including filing	Unindexed footage — findable only by people who already know it exists
Structured shadowing and reverse shadowing	Judgment-adjacent tasks; live verification that documented items are truly recoverable	Half-day sessions over weeks	Scheduled as a courtesy, cancelled under workload pressure
CMMS / IWMS record enrichment	Asset facts: feeds, dependencies, model quirks, vendor and parts specifics	Minutes at the asset with the record open	Enrichment sprints without ownership; fields decay to their prior state
Who-to-call relationship map	The relationships layer: utilities, authorities, tenant contacts, specialty vendors	1–2 hours per team	Names captured without context; map not refreshed when counterparts change
Vendor documentation clauses	Knowledge held by service-provider staff	Contract drafting time	Boilerplate obligations that no acceptance test ever verifies

Table 1. The closure-mechanisms menu. The audit's concern is that the red quadrant shrinks in systems a stranger can reach; the mechanism is matched per item.

Two principles govern the matching. First, *the record must land in a system that survives its author* — a beautifully annotated capture stored in a personal drive fails the Handoff Test as completely as no capture at all,

so every mechanism's final step is filing in the mapped system of record from Section 4, Property three. Second, *prefer the mechanism the holder will actually complete*. A veteran who will never write a procedure will happily narrate one on camera while pointing at the equipment; a technician allergic to paperwork will correct a CMMS record on the spot if the record is open on a tablet in front of the asset. Levaru's companion research treats the visual mechanisms in depth — the annotated 360-degree capture is the natural closure for items whose content is fundamentally spatial [2] — but the audit is agnostic: its concern is that the red quadrant shrinks, measurably, in systems a stranger can reach.

6. Three pillars: where the audit pays back first

The Continuity Audit earns its keep at three kinds of moment — the three named in this paper's subtitle. Each is examined in turn, because each stresses the organization differently and each converts audit outputs into a different kind of value.

6.1 Retirements and planned departures

The retirement is the canonical case, and the one in which the standard organizational response fails most visibly. The pattern is familiar enough to be a genre: the date is known a year out; the final months arrive burdened with the ordinary workload plus farewell obligations; someone remembers, six weeks from the end, to schedule "knowledge transfer sessions"; the sessions produce a binder and a folder of recordings organized by whatever occurred to people in the room; and the first winter after the departure produces the phone call — to a person now fishing, contractually absent, or simply entitled to be done — that reveals what the binder missed. Nothing in that pattern involves anyone failing at their job. It is the predictable output of treating a foreseeable, scoped, multi-year knowledge-transfer problem as a farewell-quarter task.

An operating audit changes the shape of the problem before it changes any outcome. Because the inventory exists in advance, a retirement declaration does not begin a discovery exercise; it *filters* one — the owner pulls every item held by the departing person, and the red-quadrant subset becomes the transfer plan, already scored, already scoped, with the months remaining allocated to the items that matter most rather than the ones easiest to remember. Because coverage is mapped to systems, transfer work lands where successors will look rather than in a binder whose location will itself become located knowledge. Because the Handoff Test gives the role a score, the organization can watch the number move as closure lands — and can state, on the last day, not "we did some sessions" but "the role's recoverable share moved from 40 to 85 percent, and here is the accepted-risk list for the remainder." And because elicitation happened while the holder had years rather than weeks, it happened at conversational pace, in the spaces themselves, with time to run the live test and fix what it exposed.

The same machinery serves the arrival that follows the departure. A new hire inheriting a role with a current inventory receives, in effect, a syllabus: here are the located-knowledge items this role depends on, here is where each is documented, here are the accepted gaps, and here is the veteran time you should spend on judgment rather than on 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 [4] — and while a facility team should measure its own ramp rather than adopt anyone else's, the mechanism is not in

doubt: the slowest weeks of a new technician's first year are spent rediscovering, one work order at a time, exactly the layer the audit exists to have already captured.

6.2 Contract changes and vendor transitions

If the retirement is the canonical case, the contract transition is the industrial-scale one. With outsourced delivery carrying a majority of a market measured in the trillions of dollars [8], the people holding a building's located knowledge are, in a great many portfolios, employees of a service provider — and a changeover replaces not one holder but a shift's worth of them, on a single date, across every trade the contract covers. The structural problem is misaligned incentive at the worst possible moment: the outgoing provider is demobilizing, reassigning its best people to accounts it is keeping, and has no commercial reason to fund careful transfer beyond what its contract compels; the incoming provider is mobilizing, does not yet know the buildings well enough to know what to ask; and the owner, who alone spans the transition, typically discovers the located-knowledge gap only when the first season stresses the systems.

The audit gives the owner three levers, corresponding to the three phases of a contract's life. **At procurement**, the inventory becomes contract language: because the owner knows, item by item, what located knowledge the operation depends on, transition-out obligations can be written specifically — not "provider shall cooperate in orderly handover" but a schedule of named knowledge items whose documented closure, verified against the owner's systems of record, is a condition of final payment. The construction-handover world learned this lesson a generation ago and encoded it in structured deliverables such as COBie [17]; the audit extends the same logic from commissioning data to operating knowledge. **During the contract**, cadence does the work: an annual inventory refresh that includes provider staff keeps the owner's coverage map current and keeps the closure obligation from ballooning into a final-quarter crisis — with the incidental benefit that knowledge documented in the *owner's* systems throughout the term is knowledge the owner actually owns, whatever happens commercially. **At transition**, the inventory is the checklist: incoming staff are onboarded against it exactly as a new hire is in Section 6.1, the live Handoff Test becomes an acceptance procedure for the handover, and the accepted-risk register tells the new provider precisely which gaps they are inheriting by design.

The insourcing decision and the provider consolidation are the same case with the roles rearranged, and the audit serves them identically. What all the variants share is that the transition date is contractual — known to the day, often years ahead — which makes the located-knowledge loss it threatens the single most schedulable risk in facility operations, and the least excuse-worthy to absorb unprepared.

6.3 Reorganizations, role changes and unplanned absence

The third pillar is the quiet one. No one leaves the payroll when a portfolio is consolidated, a region is restructured, maintenance is centralized, or a hands-on supervisor is promoted into a planning role — and so no alarm fires, no transition plan is commissioned, and the located knowledge involved is assumed, by default, to have moved with the org chart. It has not. The veteran now covering the north campus still holds the south campus in their head, but holding is not availability: the fifteen-minute walk-over becomes a scheduling negotiation, the radio call becomes a favor drawn on a depleting account, and within a year the knowledge is, operationally, gone — its holder alive, employed and two buildings away. Reorganizations also break the *relationships* layer of located knowledge, which inventories persistently underrate: who at the utility answers after hours, which tenant contact can authorize access, which inspector wants to see what — items whose holder was never the role but the person, and whose loss is discovered one awkward phone call at a time.

An operating audit treats a reorganization the way Section 4's triggers intend: as an event that fires a targeted refresh *during planning*, while moves can still be sequenced. The mechanics are the retirement machinery applied laterally — pull the items held by everyone whose role is changing, re-score coverage in light of where they are going, and schedule closure for the red-quadrant residue before the effective date rather than after the first failure. Organizations that run the audit through a restructuring frequently find it altering the restructuring itself in small, cheap ways: a move delayed one quarter to complete a closure sprint, a dual-coverage overlap added for a critical building, a promotion paired with a standing half-day per month in the old spaces until the score recovers.

The limiting case of this pillar — and the argument that closes it — is the absence nobody schedules. Illness, injury, family emergency and simple resignation-with-two-weeks produce, without warning, exactly the situation the Handoff Test simulates: the holder unreachable, the systems on their own. Every mechanism in this paper is, in the end, preparation for that morning. A team that has inventoried its located knowledge, closed its red quadrant and verified the closure live does not experience the unplanned absence as an operational emergency layered on a human one; it experiences it as the contingency it planned for — which is, reduced to a sentence, what continuity means.

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 [2][16]: 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 team, one building cluster, or the three most-consulted people in the portfolio — the two-week call tally of Section 5.2 will identify them if opinion cannot. Establish governance in its minimal viable form before the first interview: a one-page charter naming the owner, the rules for handling people-adjacent information, who sees which outputs, and the performance-evaluation firewall of Section 10, communicated to every participant in plain language. Baseline the metrics of Section 9 from day one, before any closure work, because the pilot's verdict depends on the before-and-after. Then run the instruments in order — elicitation, inventory, scoring — and spend the pilot's closure budget exclusively on the red quadrant, using whichever mechanisms from Section 5.4 the holders will actually complete. End the pilot with the live Handoff Test on at least one role, and report three numbers: the baseline, the movement, and the cost in hours. A pilot that cannot beat its baseline is treated as a finding, not a failure — it will usually mean the scoring flattered the coverage, which is itself worth knowing before scaling.

Walk — portfolio extension under written cadence. Extend the inventory across roles and buildings in criticality order, at whatever pace the closure capacity can honestly sustain — an inventory that grows faster than its red quadrant shrinks is accumulating anxiety, not continuity. Put the cadence in writing: a standing annual refresh, plus the event triggers of Section 4 wired to the places the events first become visible — resignation notice to the audit owner as a standard HR routing, contract final-year flag as a standard procurement routing, reorganization planning as a standard seat at the table. Connect coverage mapping to the systems of record so that closure lands in the CMMS, the drawing set, the visual register [2] and the procedure library rather than beside them, and begin managing the closure backlog with the same weekly discipline as the maintenance backlog it resembles — because that is what it is: deferred maintenance on the organization's knowledge.

“An inventory that grows faster than its red quadrant shrinks is accumulating anxiety, not continuity.”

Run — embedded operations. The audit stops being an initiative when its verbs appear inside other documents' sentences. Closure deliverables appear in vendor contracts as transition-out schedules and in project closeout as an operating-knowledge handover alongside the as-builts. The inventory syllabus appears in

onboarding curricula. The refresh appears in the annual operating calendar next to the condition assessment, and its hours appear in the operating budget as routine rather than initiative funding. The Handoff Test score appears wherever the organization already reviews operational risk. At this stage the audit also passes its own test: its method documented, its owner's successor named, its history versioned — an audit that would not survive its founder is a finding about the audit.

8. A maturity model for knowledge continuity

Teams need a way to locate themselves before they can plan a route, and a five-level model — parallel in spirit to the documentation maturity ladder of this series' earlier papers [2][16] — serves the purpose. The levels describe the *practice*, not the people: a Level 0 organization can be staffed entirely with excellent veterans, which is in fact the most common way to be Level 0, since excellence on the payroll is what makes the absence of a practice painless right up until it is catastrophic.




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

Level	Name	What it looks like	Characteristic exposure
0	Undocumented	Located knowledge lives in people; departures are absorbed by heroics and hindsight; no inventory exists.	Every departure is a discovery exercise; unplanned absence is an operational emergency.
1	Anecdotal	Risk is acknowledged; exit interviews and farewell "brain dumps" happen; artifacts are unindexed and unowned.	Transfer depends on what the departing person recalls in their busiest weeks.
2	Inventoried	A scored inventory exists for critical roles; the red quadrant is known; a pilot has beaten its baseline.	Inventory decays without cadence; closure is opportunistic.
3	Managed	Written cadence and event triggers operate; the closure backlog is worked and measured; Handoff Test scores are tracked by role.	Program depends on its owner; not yet embedded in contracts and closeout.
4	Embedded	Closure is written into vendor contracts, project closeout, onboarding and the operating calendar; the audit survives its founder.	Residual: judgment and craft still transfer only through mentoring time.

Table 2. A five-level maturity model for knowledge continuity. 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 done an audit* to *operating one* — cadence, triggers and backlog discipline — and it is sustained by exactly the ownership and governance that Properties two and five of Section 4 insist on. Level 4 is not a destination so much as a disappearance: the audit stops being visible as a program because its components have been absorbed into contracts, closeouts, onboarding and the operating calendar, which is where continuity 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 closure 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 audit is repaying its effort. Every metric below can be captured with existing systems or a two-week sampling exercise, and each should be baselined from day one of the crawl stage — afterwards is too late, because closure work contaminates the baseline it was supposed to be measured against.

“The correct benchmark is the team's own baseline — a pilot that cannot beat it has produced a finding, not a failure.”

Metric	What it measures	How to baseline it
Handoff Test score (by role)	Coverage-weighted share of a role's located knowledge recoverable from systems by a competent stranger	Score the role's inventory items; confirm with one observed live test
Concentration index	Share of sampled located-knowledge queries resolved by the three most-consulted people	Two-week self-kept tally of questions arriving by radio, phone and doorway
Time to independent competence	Weeks until a new hire completes the role's ten most common tasks unassisted	From recent hires' records; then per hire going forward
Rework attributable to missing located knowledge	Repeat visits, delays and reassignment chains caused by information that existed but could not be reached	Two-week work-order sample, coded by a coordinator
Closure backlog burn-down	Red-quadrant items closed per quarter, against items newly added	From the audit's own backlog once scoring is complete

Metric	What it measures	How to baseline it
Inventory recency	Share of inventory items reviewed within their cadence window	From the inventory's review dates; reported beside every other metric

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 **concentration index** — the share of sampled located-knowledge queries resolved by the three most-consulted people — is the program's early-warning gauge: a team can hold a respectable average Handoff Test score while concentrating its genuinely critical knowledge in two heads, and the index catches what the average conceals. It is also the metric most likely to be misread as being *about* those two people, which is why Section 10's firewall matters and why the index is reported at team level, never as a league table of the consulted. **Inventory recency** is the metric that keeps the others honest over time: closure percentages computed against a stale inventory measure progress against a team that no longer exists, so the share of the inventory reviewed within its cadence window is reported alongside 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-person technical team across a 25-building portfolio, three departures anticipated over 24 months, running the crawl and walk stages in one year.

An illustrative year-one model — every input a stated assumption	Hours
Assumed case: 20 technical staff; 25-building portfolio; three departures anticipated over 24 months; crawl and walk stages completed in year one.	—
Effort — audit design, charter and governance setup	16
Effort — elicitation walk-throughs (20 staff at ~2.5 h each)	50
Effort — facilitation, notes and inventory assembly	30
Effort — scoring workshops and closure planning	40
Effort — red-quadrant closure (~40 items at 4–6 h, mixed mechanisms)	160–240
Year-one effort, total	~300–390
Recovered — interrupt load redirected to systems (5 located-knowledge queries per working day at ~0.5 h combined asker-and-answerer cost; one-third redirected)	~200
Recovered — new-hire ramp reduction (2 hires; 3 weeks shorter at a 50% effectiveness gap)	~120
Recovered — repeat visits avoided (1 per building per quarter at 1.5 h; half eliminated)	~75
Recovered — one planned departure executed against the inventory rather than a farewell-quarter scramble (conservative rediscovery avoidance)	~150
Year-one recovery, before any avoided incident, dispute or failed handover	~545

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.

On these assumptions, the first year spends roughly 300 to 390 hours of effort and recovers on the order of 500 to 600 labor-hours — before counting a single avoided emergency escalation, disputed contract handover, failed audit finding or retirement that lands without a plan, which is where the sums that dwarf this table actually live. The model's point is the shape, not the outputs: audit and closure effort is *bounded and predictable* — it scales with team size and shrinks after year one to cadence maintenance — while the costs it offsets are *open-ended and compounding*, growing with every departure, transition and reorganization 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 coverage is better than it feared, or its scoring flattered it — and one live Handoff Test will say which.

10. Governance: auditing knowledge, not people

A Continuity Audit gathers information from individuals about the organization's dependence on them, and produces artifacts that map where that dependence concentrates. Handled carelessly, the same exercise that protects the organization can corrode the trust it depends on — a technician who suspects the inventory is a redundancy-planning instrument, or a pre-negotiation dossier on their value, will contribute exactly the inventory such an instrument deserves. Governance is therefore not a compliance appendix to the audit; it is a condition of getting true answers, and this paper treats it as non-negotiable in the same spirit as its predecessor treated the governance of building imagery [2].

“The first exception ends the program's supply of honesty — permanently.”

Five rules constitute the minimum. **First, the firewall:** audit material feeds no performance evaluation, no compensation decision and no workforce-reduction analysis — stated in the charter, said aloud in every interview, and kept, because the first exception ends the program's supply of honesty permanently. The audit measures the state of the organization's systems, and an inventory full of undocumented items is a finding about the systems, never about the people who carried what the systems failed to. **Second, participation framed as legacy, not extraction.** The honest pitch — that the exercise ensures a career's accumulated knowledge outlives the accidents of staffing, and spares its holder the 2 a.m. calls of a retirement spent as the organization's unpaid

memory — is also the accurate one, and veterans hear the difference. **Third, access control on outputs.** The scored inventory and the concentration analysis are sensitive operational records: they describe single points of failure by name. Role-based access, a named-persons distribution for anything that identifies holders, and revocation wired into offboarding apply to the audit's artifacts exactly as they apply to the building imagery and credentials the artifacts reference. **Fourth, data-protection compliance by design.** Interview notes, call tallies and holder fields are personal data within the meaning of regimes such as the EU's General Data Protection Regulation [14]; collection is minimized to operational necessity, purposes are stated, retention is rule-based and deliberate rather than indefinite by neglect, and where works councils or collective agreements govern workplace monitoring, the audit is designed with them rather than around them. **Fifth, dignity in the mechanics.** Interviews are scheduled inside working hours at working pace; the live Handoff Test is framed and debriefed as a systems test with the role's holder as its expert consultant, not its examinee; and closure work assigned to a veteran is workload, planned and credited as such. None of these rules is costly. All of them are cheaper than running the audit twice because the first run was answered defensively.

11. Standards and frameworks alignment

The Continuity Audit 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 audit work inside obligations they already carry.

Framework	Where the audit fits
ISO 30401:2018 — Knowledge management systems [9]	Supplies a facility-native method for determining necessary knowledge, addressing knowledge-loss risk and operating knowledge as a managed resource under cadence and ownership.
ISO 41001:2018 — Facility management systems [10]	Operationalizes the support and planning clauses on competence and documented information at the plant-room layer; audit artifacts serve as objective evidence.
ISO 55000 series — Asset management [11]	Extends asset-information requirements to the operating knowledge that determines asset performance; aligns the closure backlog with asset-risk registers.
IFMA competency framework [15]	Maps to the leadership and strategy, facility information management and technology management, risk management, and performance and quality competencies.
COBie / BIM handover conventions [17]	Extends structured-handover logic from commissioning data to operating knowledge; transition-out schedules mirror COBie's deliverable discipline.
GDPR and data-protection regimes [14]	Governs the people-adjacent data the audit gathers: minimization, stated purpose, rule-based retention and access control.

Table 5. Standards and frameworks alignment.

The alignment with ISO 30401 deserves the emphasis. The knowledge-management systems standard asks organizations to determine the knowledge necessary for their operations, address knowledge-loss risk, and treat knowledge as a managed resource under leadership commitment and continual improvement [9] — requirements that most facility organizations, if they encounter the standard at all, experience as abstract. The

Continuity Audit is a concrete, facility-native way to meet them: the inventory determines the necessary knowledge, the matrix and closure backlog address the loss risk, and the cadence, ownership and metrics supply the management-system machinery. An organization pursuing ISO 41001 certification, or already holding it, will find the audit slots into its support and planning clauses the same way — as the missing operational layer between the requirement to ensure competence and documented information, and the plant-room reality those requirements are ultimately about [10].

12. Limitations and outlook

Limitations. This paper's claims are bounded, and the boundaries deserve the same prominence as the argument. *The audit inventories the locatable layer, not judgment.* Diagnostic intuition, vendor-negotiation craft and the feel for a chiller's sound the week before it fails are developed, not documented; the audit's contribution to them is indirect — it frees mentoring time from address-transfer so the time can go to judgment-transfer — and any program promising to "capture everything a veteran knows" is overpromising by roughly the amount that matters most. *Scoring is structured judgment, not measurement.* Two facilitators can score the same item differently; the defenses are coarse scales, paired scoring for critical items and the live test as arbiter, and the residual subjectivity is a reason to track movement rather than worship absolute values. *The inventory decays at the rate of change,* and a lapsed audit is worse than none in one specific way: it manufactures false confidence, and in a dispute or an audit of the audit, a stale inventory presented as current is a liability, not an asset — the exact failure mode this series documented for lapsed visual registers [2]. *Governance failure is program failure.* One misuse of audit material in a personnel decision does not damage the program; it ends it, whatever continues under its name. And *the audit is not for everyone in equal measure:* a five-person team in a single small building, with genuine redundancy of knowledge across all five, may honestly conclude after a one-day exercise that its red quadrant is empty — the conclusion is legitimate, and the one-day exercise that produced it was the audit, at exactly the scale the risk deserved.

Outlook. Three developments will change the audit's mechanics without changing its logic. Speech-to-text and language models are already collapsing the cost of elicitation's most expensive step — turning a narrated walk-through into structured, searchable inventory entries — and will make the two-week call tally continuous and automatic where teams choose to instrument their channels, under governance that will matter more, not less, as the instrumentation improves. Text mining over CMMS history will surface located-knowledge hotspots — the reassignment chains and empty closure notes of Section 5.2 — as a standing dashboard rather than a periodic archaeology. And the convergence this series has traced elsewhere [16] will pull the inventory toward the building's spatial record, so that the knowledge item, its documentation and the place they describe resolve to the same address. Every one of these developments makes capture cheaper and none of them supplies cadence, ownership, honesty or a firewall — 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 is not awareness but instrumentation: a way to see the exposure before the departure, compare it across roles and buildings, spend closure effort where it matters, and prove afterwards that the effort worked. The Continuity Audit is that instrumentation — an inventory that makes located knowledge visible, a matrix that makes it comparable, a closure backlog that makes it workable, a test that makes the result verifiable, and the governance that makes the whole exercise answerable with the truth.

The practice asks little that facility teams do not already know how to do. It is, in structure, preventive maintenance — inspection on a cadence, findings in a backlog, work orders with owners and dates — applied to an asset class the CMMS has no module for and the balance sheet has no line for, which nonetheless determines how every other asset performs the morning after a veteran's last shift. The retirements are scheduled. The contract end dates are signed. The reorganizations are already on someone's whiteboard. The only genuinely open question is whether the knowledge those events will test has been inventoried, closed and verified — or whether the organization will meet its most foreseeable risk the traditional way, one radio call after it stops being answered.

References

- [1] 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 [2].
- [2] Hanov, N. *The Visual Asset Register: How 360-Degree Documentation Improves Work Orders, Inspections and Staff Handoffs*. Levaru, January 2026. levaru.co/white-papers/visual-asset-register/
- [3] 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 [2].
- [4] 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.
- [5] 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.
- [6] 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 [2].
- [7] 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.
- [8] 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.
- [9] ISO 30401:2018. *Knowledge management systems — Requirements*. International Organization for Standardization, Geneva.
- [10] ISO 41001:2018. *Facility management — Management systems — Requirements with guidance for use*. International Organization for Standardization, Geneva.
- [11] ISO 55000 series. *Asset management*. International Organization for Standardization, Geneva.
- [12] DeLong, D. W. *Lost Knowledge: Confronting the Threat of an Aging Workforce*. Oxford University Press, 2004.
- [13] Polanyi, M. *The Tacit Dimension*. University of Chicago Press, 1966.
- [14] Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation).
- [15] International Facility Management Association. IFMA core competencies of facility management (competency framework).
- [16] 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/
- [17] 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 · January 2026

Minimum Viable Digital Twin for Existing Buildings · 2026

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