LEVARU WHITE PAPER

Minimum Viable Digital Twin for Existing Buildings

A Facility Manager's Practical Framework for 360° Reality Capture, Structured Asset Data and Operational Integration

Nikolai Hanov

Published by Levaru · June 2026

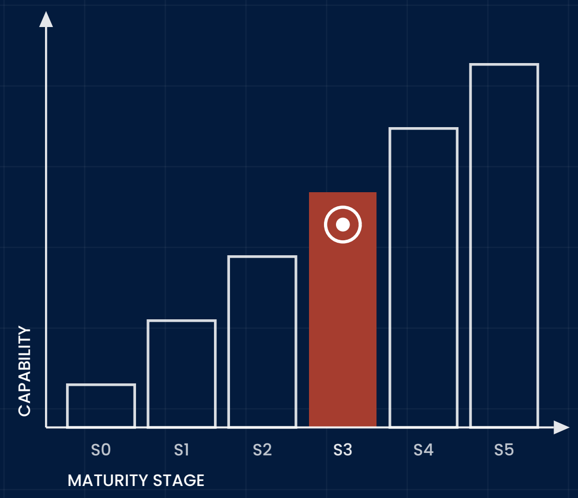

Abstract

Digital-twin guidance in the built environment is written largely from the perspective of new construction, where a coordinated model, a structured handover and a project information framework already exist. Facility managers responsible for buildings that are twenty, forty or eighty years old start elsewhere: incomplete drawings, an asset register of uncertain accuracy, renovations that never reached the record set, and operational knowledge held by two or three long-serving technicians. For those buildings, the standard advice — model the asset, instrument it, then analyse it — describes a destination without a route. This paper defines a smaller and more defensible starting point. A **minimum viable digital twin** is the least amount of connected spatial, asset and workflow information that measurably improves one operational outcome, is synchronized with the physical facility at a stated frequency and fidelity, and can be governed and audited. The framework sets out six minimum requirements — a defined operational outcome; navigable spatial context; unique location and asset identities; connections to authoritative information and workflows; defined synchronization and change management; and governance, cybersecurity and trust — and supports them with a six-stage maturity ladder, a vendor-neutral comparison of eight capture methods, a minimum asset dataset, a fourteen-step implementation sequence, a pilot roadmap bounded to one building and three use cases, transparent business-case formulas, and governance and procurement guidance covering data ownership, export rights and continuity after vendor termination. The distinction that matters most concerns 360° imagery: a virtual tour is a reality-capture layer of real and immediate value, and it becomes part of a minimum viable digital twin only when identity, connection and synchronization are added. The recommendation is narrow by design: begin with one operational problem, one governed spatial baseline and one trustworthy set of linked asset information, and add geometry, sensors, analytics and automation only when a demonstrated operational need justifies them.

KEYWORDS

digital twin · existing buildings · 360° reality capture · facility management · asset data · CMMS · synchronization · maturity model · ISO 19650 · COBie

Suggested citation: Hanov, N. (2026). *Minimum Viable Digital Twin for Existing Buildings: A Facility Manager's Practical Framework for 360° Reality Capture, Structured Asset Data and Operational Integration.* Levaru. levaru.co/white-papers/minimum-viable-digital-twin/

About this series: This is the second paper in Levaru's research series on the operational record of buildings, following *The Visual Asset Register* (January 2026).

Vendor neutrality: This paper names no commercial platform, endorses no product or service, and describes capability categories rather than implementations — capability requirements survive product cycles, and product names do not. It was prepared by Nikolai Hanov and published by Levaru, an organization active in the built-environment information field; readers should weigh the guidance with that affiliation in mind and test every recommendation against their own operating conditions.

Contents

1. Existing buildings have a different starting point

Guidance on digital twins in the built environment generally assumes a coordinated model exists. A position paper from the National Institute of Building Sciences (NIBS) treats building information models as the detailed, relatively static representation that a digital twin can animate and operationalize, with the model serving as the foundational structure on which twins are layered [1]. That logic is sound where a model has been produced and handed over — a condition most operating buildings do not meet.

The starting condition is usually some combination of the following. Record drawings are incomplete, superseded, or accurate only for the original construction. Equipment photographs sit in personal phone galleries and shared drives with no link to the asset they depict. The asset register holds entries migrated from a spreadsheet, duplicated during a system upgrade, or created in a hurry. Renovations — a chiller replacement, a tenant fit-out, a fire-alarm panel upgrade — were executed competently but never reached the record set. Work-order history is rich in text and poor in context: it says what was done without showing where, and assumes a reader who already knows the building.

Most consequentially, much of the building's operational knowledge is not written down at all. It is held by individual technicians who know which isolation valve actually serves the east riser, which access hatch sits behind the stored furniture, and which panel schedule is wrong. That knowledge leaves when they do.

These conditions are not local failures. Research published by the IFMA IT Community with Autodesk reports that 85 percent of surveyed facility operators struggle with poor-quality data, that 80 percent of firms do not consistently involve operations teams in design and construction, and that half consider their facility management teams to lack sufficient building information modelling (BIM) knowledge [2]. The same research reports that most data generated during construction is never used downstream, and that more than half of facility teams surveyed waited over six months after project completion for final closeout documentation [3]. Two decades earlier, a study prepared for the U.S. National Institute of Standards and Technology estimated the annual cost of inadequate interoperability in the U.S. capital facilities industry at US$15.8 billion, of which US$10.6 billion — roughly two-thirds — fell on owners and operators rather than designers or constructors [4]. The cost of disconnected information has always landed disproportionately on the people who run buildings.

> *"The cost of disconnected information has always landed disproportionately on the people who run buildings."*

Three assumptions imported from new construction cause most of the difficulty.

The completeness assumption. New-construction delivery aims at a comprehensive dataset covering every specified system. Applied retroactively it becomes a survey of the whole building before any operational benefit is available, and such programmes tend to exhaust their budget during data collection.

The geometry-first assumption. New-construction workflows produce geometry first and attach data to it. In an operating building, geometry is the expensive part and often the least urgent: a technician dispatched to an unfamiliar mechanical room needs to know which unit, where it is, what it serves and what was done to it last — not its exact dimensions.

The single-handover assumption. New construction has a defined moment when information transfers to the operator. An existing building has none. Information arrives unevenly from contractors, service vendors, inspections and internal projects, so a programme that treats collection as a one-time event produces a record that begins ageing immediately.

The practical consequence is that existing buildings should begin with an operational use case rather than a technology. The NIBS subcommittee makes the same point, describing use cases as the narrative preceding technology selection and the requirement against which implementation is tested [5]. Stated plainly: decide which recurring problem is worth solving, determine the smallest set of information that would solve it, and capture that. Scope discipline is not a compromise here; it is what makes the work finishable.

THE STARTING CONDITIONS, IN PUBLISHED FIGURES

85%	80%	6+ mo	$15.8B
of surveyed facility operators struggle with poor-quality building data [2]	of firms do not consistently involve operations teams in design and construction [2]	waited by more than half of facility teams for final closeout documentation [3]	annual cost of inadequate interoperability in U.S. capital facilities — two-thirds borne by owners and operators [4]

Sources: IFMA IT Community & Autodesk [2][3]; NIST GCR 04-867 [4]. Full citations in the References.

> **Scope of this paper.** The framework addresses existing buildings in operation, not new-construction information delivery, which is governed by established project information-management practice. The paper is not a standard: it interprets published standards and guidance for practitioners and does not substitute for them. Financial guidance in Section 9 supplies measurement formulas, not benchmarks — no savings percentage in this paper is offered as evidence of value in any particular building. And statements about capture methods describe general characteristics of method categories, not the performance of any specific device or service.

2. What a digital twin is — and is not

Terminology here is unsettled, and the confusion has operational consequences: organizations procure one thing, expect another, and cannot tell whether they received either. The Digital Twin Consortium defines a digital twin as an integrated, data-driven virtual representation of real-world entities and processes, with synchronized interaction at a specified frequency and fidelity, motivated by outcomes and driven by use cases [6]. The NIBS Digital Twin Integration Subcommittee adopts that definition and adds one framed for the built environment: a fit-for-purpose and intelligent virtual representation of an asset, synchronized at specific frequencies, with an existing or planned connection between the virtual and physical twin that may include analysis and the ability to actuate physical changes [7].

> **Definition.** A **minimum viable digital twin** is the least amount of connected spatial, asset and workflow information that measurably improves one operational outcome, is synchronized with the physical facility at a stated frequency and fidelity, and can be governed and audited.

Two elements do the real work. **Synchronization** separates a twin from a record; **specified frequency and fidelity** makes the claim testable. A representation synchronized quarterly at room-level fidelity is a legitimate digital twin for use cases quarterly room-level information can serve, and is not one where hourly equipment-level state is required. The question is never whether something is a twin in the abstract, but whether its frequency and fidelity match the decision it supports.

> *"Synchronization separates a twin from a record; specified frequency and fidelity make the claim testable."*

2.1 Distinguishing the terms

Term	Working definition	What it does not do
Digital archive	A stored collection of documents, drawings and images relating to a facility.	Provides no spatial navigation and no defined relationship to current physical conditions.
360° virtual tour	A navigable set of panoramic images allowing a user to move through captured positions in a facility.	Carries no structured asset data, no link to operational records, and no synchronization commitment.
Visual facility record	A virtual tour that has been organized, dated, labelled by location and retained as a controlled record.	Still does not connect to asset identities or workflows; it documents appearance, not operational state.
Digital model	Any digital representation of the facility, including CAD drawings and geometric models.	Has no automatic or manual data flow from the physical asset.
Building information model	A shared digital representation of the physical and functional characteristics of a built asset.	Is typically produced through a project and updated manually; on its own it represents design or as-built intent, not current condition.
Digital shadow	A virtual representation with data flowing from the physical asset to the digital representation only.	Does not return instructions or changes to the physical asset.
Digital thread	The interconnected mechanism that links information across systems, lifecycle stages and time.	Is not itself a representation of the facility; it is the connective tissue between representations and records.
Digital twin	A virtual representation synchronized with the physical asset at a specified frequency and fidelity, serving a defined outcome.	Does not require simulation, sensors or automation in every case.

Term	Working definition	What it does not do
Digital-twin system	A system of systems that implements one or more digital twins across an organization.	Is an operating capability, not a single software purchase.

Table 1. Terms that are frequently confused. Definitions are adapted from the Digital Twin Consortium and the NIBS Digital Twin Integration Subcommittee where those bodies have published definitions [6][7]; working definitions are supplied where they have not.

2.2 Synchronization, frequency and fidelity

Synchronization describes how, and how often, the representation is brought back into agreement with the facility. Three properties should be stated for every element of a twin.

Frequency is how often the element is refreshed: continuously, daily, monthly, annually, or on a trigger event such as equipment replacement or completion of a renovation. Trigger-based synchronization often suits existing buildings better than calendar-based, because buildings change episodically. Operational-phase information management follows the same logic, treating trigger events as the point at which asset information is updated [8].

Fidelity is how precisely the representation reflects the asset: whether location resolves to a building, floor, room or coordinate; whether a record identifies a category, a type or a specific serialized unit; whether imagery is current to the month or the year.

Direction is whether information flows one way or both. A representation receiving condition data but sending nothing back is a digital shadow; one that can also cause physical change — a control action, a work order dispatching a technician — closes the loop. Most minimum viable twins in existing buildings receive condition information and act on it only through human-mediated workflow.

Updates may be manual or automated, and manual updating is legitimate provided it is scheduled, assigned to a named owner and evidenced. A quarterly walk-through that verifies nameplate data and recaptures changed spaces is a synchronization process. An intention to update the model "when things change" is not.

2.3 Why most use cases do not require continuous data

Continuous data is expensive, and its value depends on whether a decision changes as a result. Remote work-order triage needs accurate location, correct equipment identity and a current visual record — none of which changes minute to minute. Vendor scoping needs access routes and clearances; onboarding needs spatial familiarity; emergency planning needs isolation points and egress. For these, monthly to annual synchronization with trigger-based updates is sufficient, and describing such a twin as operating in real time would be inaccurate.

Continuous data earns its cost where a decision is genuinely time-sensitive: temperature excursion in a laboratory or cold store, differential pressure in a containment space, runtime accumulation driving condition-based maintenance. Those are reasons to instrument specific equipment, not a building.

The same restraint applies to simulation and automation, neither of which a twin requires. Analytical and predictive capability sits at the upper end of the maturity ladder and depends on data most existing buildings do not yet have; building the connected, governed base first is what makes it possible later.

2.4 Stating the claim precisely

The practical discipline that follows from these definitions is a short written statement for any representation the organization operates or procures: what is represented, at what fidelity, refreshed at what frequency or on which triggers, connected to which systems of record, and owned by whom. One sentence per element is enough. A representation described as “panoramas of all mechanical spaces, refreshed annually and on project closeout, linked to the maintenance system, owned by the facility information manager” can be tested, budgeted and audited. One described as “a digital twin of Building C” cannot.

The statement also settles disputes before they occur: when a stakeholder expects live equipment state from a representation whose stated frequency is annual, the gap is visible in a sentence rather than discovered in an incident review. The six requirements in Section 3 give the statement its content, and the maturity ladder in Section 4 gives it its vocabulary.

> **A 360° virtual tour alone is not a digital twin.** It is a visual or reality-capture layer. It becomes part of a minimum viable digital twin when connected to structured operational information and synchronized with the physical facility at a defined frequency and fidelity. The distinction is not academic: it determines what the organization can rely on the representation to tell it, and what must still be verified on site.

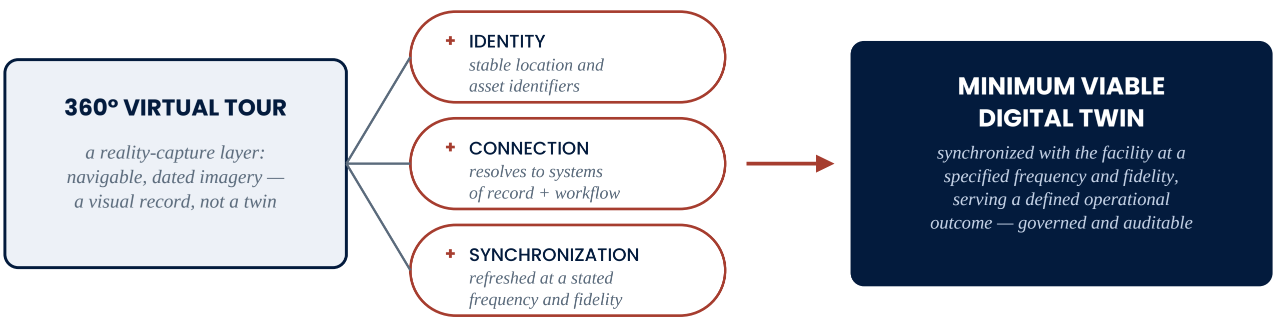

Figure 1. What a 360° virtual tour provides, and what must be added for it to form part of a minimum viable digital twin.

3. Six minimum requirements

A representation qualifies as a minimum viable digital twin when it satisfies six requirements. Each is stated with its definition, the operational reason it exists, the least implementation that satisfies it, the failure most often observed, an example, and the evidence an auditor or incoming manager could inspect.

3.1 A defined operational outcome

Definition. A named, recurring operational problem the twin exists to improve, expressed with a measure and a baseline.

Operational reason. Without one there is no basis for deciding what to capture, what to omit, or when the work is finished, and scope expands until the budget ends. Use cases are the requirement against which implementation is tested, not a justification written afterwards [5].

Minimum implementation. One outcome statement per use case naming the problem, the affected role, the measure, the baseline, the target and the accountable owner. One page is sufficient.

Common failure. The outcome is written as a technology objective — “create a digital twin of Building C” — rather than an operational one. It cannot be scoped or measured, and so cannot visibly succeed.

Example. A campus team finds one in five mechanical work orders needs a preliminary visit purely to identify the unit; the outcome statement targets that visit, measured from work-order records.

Evidence. A signed outcome statement per use case, with a baseline measured before implementation rather than estimated after it.

3.2 Navigable spatial context

Definition. A spatial representation that lets a user who does not know the building locate a space or asset and understand its surroundings.

Operational reason. Maintenance records describe actions, not places. A technician, vendor estimator or remote duty manager needs the approach, the clearance, the adjacent equipment and the access route. Spatial context converts a record into an actionable instruction.

Minimum implementation. Dated 360° panoramas covering in-scope spaces under a documented building, floor and space hierarchy, with visual continuity between adjacent positions. Floor plans may be raster images provided each position is indexed to a space identifier. Geometric models are not required.

Common failure. Positions are spaced too far apart, so users lose orientation and abandon the tour; or coverage is complete for occupied areas and absent for the mechanical, electrical and roof spaces where the value lies.

Example. A contractor quoting a pump replacement establishes the route from the loading dock, door widths, headroom and adjacent piping, and revises the quotation without a site visit.

Evidence. A capture log recording every space, its date, operator and space identifier, and a recapture list recording spaces that were inaccessible and why.

3.3 Unique location and asset identities

Definition. Every in-scope location and asset carries one persistent identifier, unique within the portfolio, used identically across every system that refers to it.

Operational reason. Identity allows a point in an image, a row in a maintenance system, a document and a physical label to be recognized as the same thing; without it, integration is guesswork. Open exchange schemas for facility handover rest on this principle [9].

Minimum implementation. A documented scheme covering portfolio, site, building, floor, space and asset; identifiers that are opaque and stable rather than descriptive; a rule against reuse after decommissioning; and physical labelling so the identifier is verifiable in the field.

Common failure. Identifiers encode attributes that later change — location, manufacturer, system designation — so a relocated or replaced asset carries a misleading identifier or is renumbered, breaking every historical link. The second is parallel numbering across the maintenance system, the labels and the drawings.

Example. A unit labelled AHU-3-02 in the maintenance system, on the equipment and in the tour hotspot resolves to one record; when it is replaced, the identifier persists and the service history stays continuous.

Evidence. A published identifier standard; a reconciliation report showing the match rate between hotspots, maintenance records and field labels; and an exception list with owners and due dates.

3.4 Connections to authoritative information and workflows

Definition. Each identified asset resolves to its record in the system authoritative for that information, and the relevant workflow can be started from that point.

Operational reason. A representation holding its own copy of asset data creates a second version that will diverge from the first. The twin should be a route to authoritative information, not a competing store. Where a system is described as authoritative it should be named per information type — maintenance history in the computerized maintenance management system (CMMS), space and occupancy in the integrated workplace management system (IWMS), live equipment state in the building management system (BMS), controlled documents in the document management system.

Minimum implementation. An information-ownership map assigning each field to exactly one authoritative system; a link from each asset to its record there; and at least one workflow, usually work-order creation or lookup, reachable from the twin.

Common failure. Attributes are typed into the twin platform because it is faster than integrating. Within a year two registers disagree and technicians trust neither.

Example. A duty manager receiving an after-hours alarm navigates to the mechanical room, selects the unit, reads its history through the link, confirms the isolation point, and raises a work order with the identifier populated.

Evidence. A published, dated ownership map; a sample of assets traceable from twin to authoritative record without manual searching; and no field in the twin that is also mastered elsewhere.

3.5 Defined synchronization and change management

Definition. A documented process stating how each element is updated, at what frequency and fidelity, by whom, and on which trigger events.

Operational reason. Synchronization is what distinguishes a twin from a record [6]. An unsynchronized representation degrades silently, continuing to look authoritative long after it stopped being accurate — which is more dangerous than having none.

Minimum implementation. A one-page schedule listing each element, its authoritative system, frequency, fidelity, triggers and named owner; a rule that closeout of any in-scope project or replacement includes a capture and data-update task; and a periodic verification walk with a documented result.

Common failure. Imagery is captured once, the programme is declared complete and nothing is scheduled. Two renovations later the tour shows rooms that no longer exist, and confidence collapses.

Example. Annual recapture for corridors and public areas; recapture on completion for spaces affected by projects above a defined value; quarterly nameplate verification for critical assets; immediate record updates on replacement.

Evidence. A dated schedule; capture logs showing updates on schedule; and closeout checklists carrying the capture and data-update requirement as a signed item.

3.6 Governance, cybersecurity and trust

Definition. Defined ownership, access control, privacy protection, security controls and retention rules covering the twin and everything it exposes.

Operational reason. A navigable visual record describes a facility's physical security: door hardware, camera positions, panel locations, server rooms, isolation points. Built-environment guidance places trustworthiness — security, openness and quality — alongside purpose and function [10], and operational-phase information management includes a security-minded approach for the same reason [11].

Minimum implementation. A named data owner; role-based access with restricted areas excluded or separately permissioned; privacy treatment for images containing people or tenant property; authentication aligned to the organization's identity standard; documented retention, backup and export rights; and a stated position on hosting jurisdiction.

Common failure. Access granted broadly during a pilot is never narrowed, and security infrastructure captured because it happened to be in the room stays visible to every contractor with a login.

Example. Communications rooms, security equipment rooms and cash-handling areas are designated restricted before capture, then excluded or held in a separately permissioned collection for a named list of staff.

Evidence. A current access matrix by role; a restricted-area register agreed before capture; evidence of privacy redaction; a tested restore from backup; and a verified data export.

Reading the six requirements together. They are not sequential stages. A programme that satisfies five and omits identity produces a record that cannot be integrated. A programme that satisfies five and omits synchronization produces a record that expires. The requirements function as a set, and the weakest determines what the twin can be trusted to do.

The requirements function as a set, not a sequence — the weakest determines what the twin can be trusted to do.

Figure 2. The six minimum requirements of a minimum viable digital twin. The set functions together; the weakest requirement determines what the twin can be trusted to do.

4. A digital-twin maturity model for existing buildings

Maturity ladders are useful when they describe capability rather than expenditure. The six stages below are defined by what the organization can rely on the representation to do, and by the evidence supporting that reliance.

Stage 0 — Fragmented facility information. Drawings, photographs, manuals, spreadsheets and maintenance records exist but sit separately, with no shared identifiers and no defined ownership, so retrieval depends on individual memory. Most operating buildings sit here — not through poor management, but as the accumulated result of decades of episodic change without an information framework.

Stage 1 — Navigable visual record. A dated, organized set of 360° imagery covers defined spaces and can be navigated by someone unfamiliar with the building, delivering immediate value for orientation, remote familiarization and condition documentation. **Stage 1 is normally not a digital twin.** It carries no structured asset data, no connection to operational systems and no synchronization commitment. Describing it as a twin creates an expectation of currency it cannot meet, and that misdescription is among the most common causes of lost credibility.

Stage 2 — Contextual facility model. Locations and in-scope assets carry unique identifiers attached to positions in the imagery, and documents and records are linked, so a user can navigate to an asset and reach its manual, drawing and maintenance history. **Stage 2 becomes a minimum viable operational twin only when paired with an appropriate and documented synchronization process.** Without it, this is a well-organized snapshot with a known expiry date and no scheduled renewal.

Stage 3 — Minimum viable operational twin. All six requirements are satisfied: at least one defined operational outcome shows measured improvement, synchronization is scheduled and evidenced, and governance and access control are in force. Most existing buildings can reach Stage 3 for a bounded scope within one budget cycle, and most operational value available from a digital twin in an existing building is available here.

Stage 4 — Analytical twin. Condition, runtime, energy or occupancy data is integrated at a defined frequency and analysed to prioritize capital renewal, detect drift in equipment performance or test space-utilization assumptions. Stage 4 depends on the data Stage 3 makes trustworthy; attempted earlier, analysis is applied to an asset register nobody believes and its outputs are disregarded.

Stage 5 — Predictive or closed-loop twin. Predictive models inform intervention before failure, or the representation participates in control through a bidirectional connection to building systems. Guidance for the built environment describes this actuating category explicitly while noting it is neither necessary nor desirable for every use case or lifecycle phase [12]. Stage 5 introduces operational-technology security obligations that should be addressed deliberately, with the information security function.

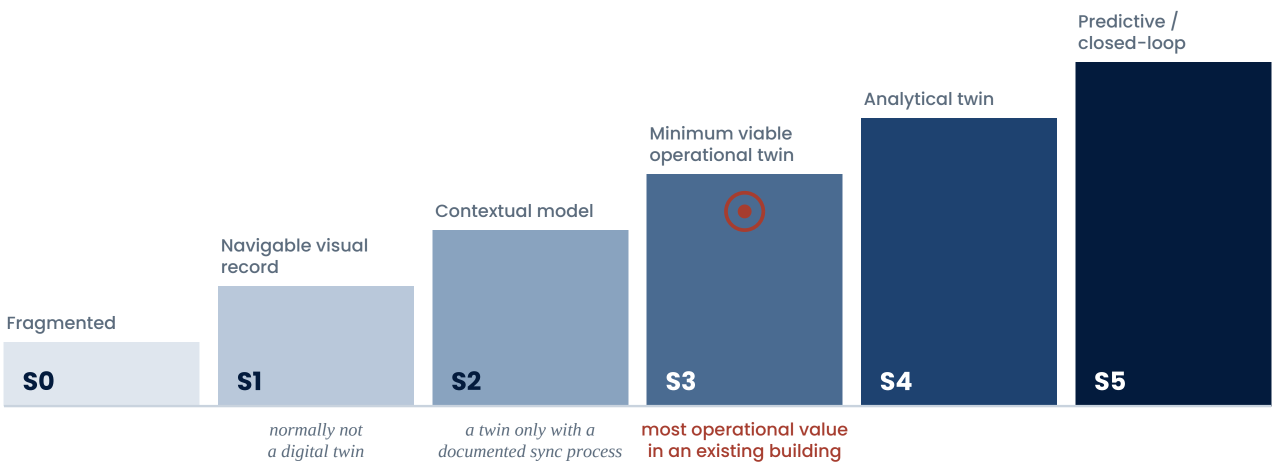

Figure 3. Six stages of digital-twin maturity for existing buildings. Stage 1 is normally not a digital twin; Stage 2 becomes a minimum viable twin only when paired with a documented synchronization process.

Stage	What can be relied on	Evidence at this stage	Typical effort profile
0 — Fragmented	Individual knowledge and manual search.	None. Retrieval time is unmeasured.	No incremental cost; recurring hidden cost in search time.
1 — Navigable visual record	What the space looked like on the capture date.	Capture log with dates and coverage list.	Days to weeks per building, depending on scope.
2 — Contextual model	Where an asset is and where its documents are.	Identifier standard; link inventory; document index.	Weeks; effort sits in identity reconciliation, not capture.
3 — Minimum viable operational twin	Measured improvement in one outcome, from information kept current on a defined schedule.	Outcome statement with baseline; synchronization schedule; access matrix; adoption data.	One budget cycle for a bounded scope.

Stage	What can be relied on	Evidence at this stage	Typical effort profile
4 — Analytical twin	Trends and comparisons informing planning and prioritization.	Data-quality reporting; documented analytical method and its limits.	Ongoing; needs data engineering and analytical skills.
5 — Predictive or closed-loop	Anticipated intervention, or control participation.	Model validation records; operational-technology security controls; change-control evidence.	Sustained programme with information security involvement.

Table 2. Maturity stages, evidence and typical effort. Effort profiles are indicative planning guidance rather than measured benchmarks, and vary with building size, complexity, access constraints and the condition of existing records.

Two rules govern movement between stages. Progression is **per use case and per scope**, not per organization: a portfolio may hold one building at Stage 3 and the rest at Stage 0, which is coherent rather than incomplete. And **stages are not skippable in substance** — an organization can buy Stage 5 technology at any time, but not the identity discipline and synchronization habit that make its outputs believable.

The ladder is also a self-assessment instrument, and the assessment is only useful when the evidence column is applied honestly. Most organizations describe themselves one stage above where their evidence places them — usually because imagery exists and is being called a twin, or because identifiers exist in the maintenance system but have never been reconciled against the field. The corrective is mechanical rather than motivational: claim the stage whose evidence can be produced on request, and record the gap to the next stage as a work item with an owner.

Used this way, the stages become shared vocabulary. A budget request framed as "moving Building C from Stage 1 to Stage 3 for three use cases" names its scope, its evidence and its finish line in one sentence, and a vendor claim framed against the same ladder can be tested rather than admired.

5. Understanding 360° reality capture

Reality capture is a family of methods with different costs, outputs and failure modes. Choosing among them becomes straightforward once the question is framed correctly: what decision must the information support, and what is the consequence of it being wrong by a measurable margin?

5.1 Comparison of methods

Method	Appropriate use cases	Complexity	Visual value	Spatial or measurement value
360° still photography	Orientation, remote triage, vendor scoping, onboarding, condition documentation, emergency familiarization.	Low. Trainable in a day; modest equipment cost.	High. Complete context in every direction from each position.	Low. Position known only to the accuracy of the floor-plan index.
360° video	Walkthrough narration, rapid route documentation, training material.	Low to moderate. Fast to capture; harder to index.	Moderate to high, but transient.	Very low. Frames are not reliably positioned.

Method	Appropriate use cases	Complexity	Visual value	Spatial or measurement value
Conventional detail photography	Nameplates, serial numbers, damage, valve tags, panel schedules, confined details.	Very low.	High for the detail; none for context.	None.
Smartphone or tablet LiDAR	Rough room dimensions, quick clearance checks, small-area scans.	Low. Available on current devices.	Moderate.	Moderate but variable; range and drift limit reliability.
Photogrammetry	Textured three-dimensional representation of objects, façades or plant rooms.	Moderate to high. Capture discipline and processing time.	High.	Moderate to high, dependent on control, overlap and texture.
Mobile mapping and SLAM	Rapid point-cloud coverage of large or circulation-heavy areas.	Moderate to high. Specialist equipment and operators.	Moderate.	Moderate to high; generally below static scanning.
Terrestrial laser scanning	Retrofit design, clash-sensitive coordination, structural and MEP as-built documentation.	High. Specialist service, planning and processing.	Moderate to high with imagery.	High. The reference method where dimensions carry cost.
Existing CAD, BIM and floor plans	Space hierarchy, room numbering, navigation base, area reporting.	Low to moderate — the effort is validation.	Low.	Variable and often unknown; may reflect design intent only.

Table 3a. Reality-capture methods: applicability and value. Characteristics describe method categories in general terms; individual devices and services vary. Abbreviations: LiDAR, light detection and ranging; SLAM, simultaneous localization and mapping; MEP, mechanical, electrical and plumbing.

Method	Limitations	Update suitability	Typical failure modes	When another method is required
360° still photography	Not measurement-grade. Occluded equipment stays occluded. Limited detail at distance.	Excellent — cheap to repeat, easy to version by date.	Scene spacing too wide; poor lighting; operator visible; missing nadir; stitching artefacts at doorways.	When a dimension will drive a purchase, fabrication or clearance decision.
360° video	Difficult to index, version and cite. Motion blur.	Poor as a controlled record; acceptable as supplementary material.	Unstable framing; no fixed reference point; no reliable link to a space identifier.	Whenever the record must be referenced, linked or compared over time.
Conventional detail photography	No spatial context; easily orphaned from its asset.	Excellent, if each image carries an asset identifier.	Unreadable nameplates; images stored without linkage; no capture date.	When context or navigation is needed.
Smartphone or tablet LiDAR	Accuracy varies with range, lighting and surface; drift accumulates.	Good for small, frequently changing areas.	Overstated confidence in derived dimensions; corridor drift; poor results on dark or glossy surfaces.	When tolerance must be specified and verified.
Photogrammetry	Struggles with featureless, reflective or transparent surfaces; scale needs control.	Moderate — reprocessing effort is significant.	Insufficient overlap; inconsistent lighting; unverified scale.	When accuracy must be independently validated.

Method	Limitations	Update suitability	Typical failure modes	When another method is required
Mobile mapping and SLAM	Accuracy below static scanning; registration drift over long runs.	Good for periodic re-baselining of large areas.	Trajectory drift; loop-closure error; poor results in long featureless corridors.	When high accuracy is needed in a confined, critical area.
Terrestrial laser scanning	Slower and costlier; occlusion still applies; large datasets to manage.	Poor for frequent refresh; suits periodic re-baselining.	Too few setups leaving shadowed areas; registration error; deliverables the facility team cannot open.	When frequent, low-cost visual refresh is the actual requirement.
Existing CAD, BIM and floor plans	May be superseded, partial, or accurate only to original construction.	Depends entirely on whether change control exists.	Treated as current without verification; room numbering that does not match signage.	Whenever current physical conditions must be confirmed.

Table 3b. Reality-capture methods: limitations, update behaviour and failure modes.

> **Ordinary 360° imagery should not be assumed to be survey-grade or measurement-grade.** Panoramic photographs are not produced under controlled geometry, and any dimension derived from them carries unquantified error. Where a decision depends on a measurement, required accuracy should be specified in procurement using a recognized framework — the U.S. Institute of Building Documentation's Level of Accuracy (LOA) Specification defines graded accuracy levels for building documentation and expresses tolerance against a stated confidence level rather than a single figure — and the deliverable validated against it [13]. Publishing measurements taken from panoramic imagery without that basis transfers unquantified risk to whoever relies on them.

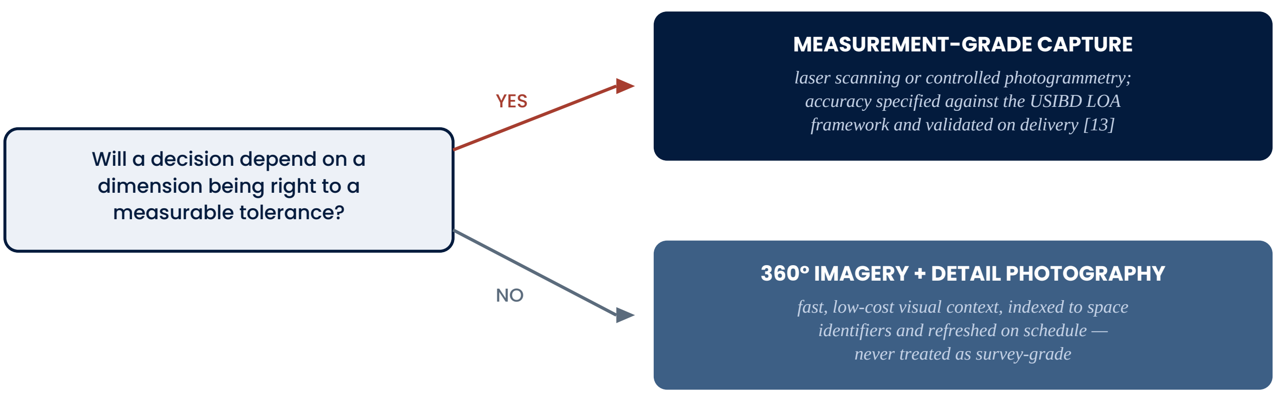

Figure 4. Selecting a capture method by decision type. Where a dimension will drive a purchase, a fabrication or a clearance decision, panoramic imagery is not sufficient.

5.2 Practical capture considerations

Capture quality decides whether the record is used or abandoned, and most of it is settled before the first image is taken. Appendix B gives the working field checklist; the reasoning behind its main provisions follows.

Scope and identifiers. Fix the capture scope in writing before mobilizing — which buildings, floors, space types and equipment rooms, and what is excluded. Every position must carry floor and room identifiers drawn from the same hierarchy the maintenance system uses, established before capture rather than reconciled afterwards.

Height, spacing and transitions. A consistent camera height at standing eye level keeps scenes coherent and makes comparison across dates meaningful. Spacing should be close enough that the next position is always visible; corridors, open-plan areas and repetitive spaces need tighter spacing than intuition suggests, because they offer few landmarks. Doorways deserve a position on each side, since one position in a doorway produces stitching artefacts and disorienting transitions. Stairs need a position at each landing; elevators are best represented by lobby positions on each floor. A visible floor or room marker in frame helps viewers confirm where they are.

Operational spaces. Mechanical, electrical and telecommunications rooms carry most of the value and the hardest conditions: tight clearances, high contrast between windows and shadowed plant, equipment obscured by equipment. Plan several positions per room, and accept that occluded items need supplementary detail photography rather than another panorama. Rooftops need positions showing access, curbs, drainage, guard-rail condition and unit orientation. Above-ceiling areas, where safe, should be captured selectively at known access points.

Lighting, reflection, movement and artefacts. Interior scenes routinely exceed the dynamic range of one exposure, so bracketed capture is needed where a window and an unlit plant room appear together. Mirrors and glazing duplicate or distort scenes and should be noted in the log. People and moving equipment cause ghosting; capture outside occupied hours where practical and record the condition where not. Stitching artefacts concentrate at scene edges, near objects and at high-contrast boundaries. Nadir and zenith coverage should be checked deliberately, and operator and tripod removal stated as a deliverable.

Privacy and restricted areas. Faces, screens, personal desks, tenant property and personally identifiable material should be blurred or excluded under a written rule agreed before capture. Restricted areas are addressed in Section 10.

Records of capture. Three records make a programme auditable: a capture log for every position with space identifier, date and operator; a recapture list for spaces missed, locked or obstructed; and image versioning retaining superseded imagery with its date. Nameplate photography taken in the same visit and attached to the asset identifier converts a visual record into the beginning of a verified asset register.

6. Minimum facility data

The temptation in an existing building is to define an ideal asset record and then fail to populate it. The alternative is a deliberately small dataset that can be completed for in-scope assets, verified and dated. Open exchange schemas for facility handover follow the same logic, organizing information around facilities, floors, spaces, systems and components [9]. The dataset below adapts that structure to what an operating facility can realistically collect in a first capture cycle.

6.1 The minimum asset dataset

Field	Purpose	System of record	Notes
Portfolio or campus ID	Groups sites under one ownership or management structure.	IWMS or property register	Required only for multi-site organizations.
Site ID	Identifies the physical site or campus.	IWMS or property register	Stable; never reused.
Building ID	Identifies the building.	IWMS or CMMS	Must match signage and lease documentation where these exist.
Floor ID	Identifies the level.	IWMS or CMMS	Follow the wayfinding convention, not the drawing convention, where they differ.
Space or room ID	Identifies the space; anchors capture positions and asset locations.	IWMS or CMMS	The most common point of failure; reconcile with door signage before capture.
Asset ID	Uniquely identifies the asset across all systems.	CMMS or EAM	Opaque and persistent; survives replacement of the physical unit.
Asset category	Classifies the asset for reporting and planning.	CMMS or EAM	Use an existing classification rather than inventing one.
System	Associates the asset with the system it serves.	CMMS or EAM	Enables impact assessment during failure and shutdown planning.
Manufacturer	Supports parts, service and warranty.	CMMS or EAM	Capture from the nameplate photograph.
Model	Supports parts, service and warranty.	CMMS or EAM	Record exactly as printed, including suffixes.
Serial number	Identifies the specific physical unit.	CMMS or EAM	Changes on replacement while the asset ID does not.
Criticality	Prioritizes maintenance, spares and capital renewal.	CMMS or EAM	Use the organization’s existing scale; do not create a parallel one.

Field	Purpose	System of record	Notes
Maintainable status	Distinguishes maintained assets from inventoried components.	CMMS or EAM	Controls whether the asset generates planned maintenance.
Current condition or status	Records observed condition at last verification.	CMMS or condition-assessment system	Always paired with the verification date.
Install or replacement date	Supports lifecycle and capital planning.	CMMS or EAM	Record as known or estimated, and label which.
Warranty information	Prevents paying for covered work.	CMMS or contract repository	Include expiry and the party holding the obligation.
Information source	Records where each value came from.	Twin platform or CMMS	Nameplate, drawing, contractor submittal, technician report, estimate.
Data-confidence level	States how far the value can be trusted.	Twin platform or CMMS	A three-level scale — verified, reported, estimated — is sufficient.
Date last verified	Makes staleness visible.	Twin platform or CMMS	The single most useful field in an inherited asset register.
Responsible data owner	Assigns accountability for accuracy.	Governance register	A named role, not a department.
Linked manuals, drawings, service records	Provides the documents needed to act.	Document management system; CMMS for service history	Store links, not copies.

Table 4. Minimum asset dataset and the system authoritative for each field. Abbreviations: CMMS, computerized maintenance management system; EAM, enterprise asset management; IWMS, integrated workplace management system.

6.2 What belongs in the twin and what does not

The division is a matter of authority rather than convenience. The twin platform should hold the spatial and visual layer — panoramic imagery and its capture metadata, positions, hotspot geometry and the mapping between a point in an image and an identifier — because no other system holds it. It should also hold links and the metadata describing their trustworthiness: information source, data-confidence level and date last verified.

Everything else should remain where it is already governed. Maintenance history, work orders, planned maintenance schedules, asset hierarchy and criticality belong to the maintenance management system; space allocation, occupancy and area data to the workplace management system or property register; live equipment state, setpoints and alarms to the building management system, read from it rather than copied; and controlled documents — manuals, drawings, certificates, commissioning records — to the document management system, where version control already exists.

The rule that prevents most downstream problems is simple: any field mastered elsewhere should be displayed by the twin, not stored by it. Where caching is required for performance, the cache should carry its source timestamp and be visibly identified as a copy. Naming the authoritative system for each field, in writing, is what lets an organization say which record wins when two disagree — and that question will arise.

> *“Any field mastered elsewhere should be displayed by the twin, not stored by it.”*

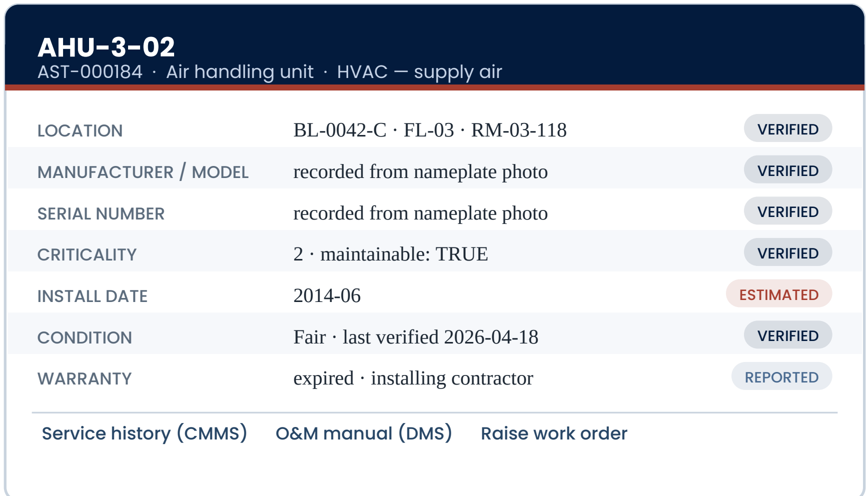

Figure 5. An example asset-information card showing the minimum dataset, with data-confidence and last-verified indicators visible to the user.

7. From capture to operations

The sequence below takes a bounded scope from decision to measured outcome. Order matters more than duration, because several steps become far more expensive if attempted out of sequence.

1. **Select use cases.** Choose no more than three recurring operational problems and write an outcome statement for each, with a measure and an owner.
2. **Establish scope.** Define the buildings, floors, space types, asset categories and systems in scope, and record what is excluded and why.
3. **Create the building and location hierarchy.** Agree portfolio, site, building, floor and space identifiers before capture, reconciled against door signage, wayfinding and the maintenance system. Conflicts resolved now cost hours; resolved later they cost a re-survey.
4. **Capture spatial information.** Execute the plan with consistent camera height and scene spacing, maintaining the capture log and recapture list throughout.
5. **Process and quality-check.** Review stitching, exposure, nadir and zenith coverage, operator removal and privacy redaction. Reject and recapture before publishing; quality problems found by users cost adoption.
6. **Identify in-scope assets.** Walk the captured spaces, list the assets in scope and photograph each nameplate.
7. **Create or validate asset identifiers.** Match each asset to an existing record or create one, recording every unmatched item on an exception list with an owner and a due date rather than resolving it silently.

8. **Link assets to records.** Establish the connection from hotspot to authoritative record, and test that it resolves from a mobile device on the network technicians actually use.
9. **Link documents.** Attach manuals, drawings, certificates and service history by reference from the document management system.
10. **Establish access controls.** Apply the role matrix, confirm restricted-area handling, and verify a test account in each role sees only what it should.
11. **Establish synchronization rules.** Publish the schedule and embed capture and data-update tasks in closeout and equipment-replacement procedures.
12. **Train users.** Train by task, not feature: triaging a work order remotely, scoping a vendor visit, verifying an asset in the field. Fifteen minutes of task-based training outperforms a platform demonstration.
13. **Measure adoption and outcomes.** Track usage by role and the step 1 measures against their baselines. Adoption without outcome improvement indicates the wrong use case; the reverse indicates a measurement error.
14. **Expand only after evaluation.** Review what was used, ignored, inaccurate and what capture actually cost, and fold those findings into the next scope before extending.

Two steps are consistently underestimated. Step 3, the location hierarchy, determines whether every subsequent link is possible, and because it is administrative rather than technical it is often skipped in the enthusiasm to begin capture. Step 7, identifier validation, routinely surfaces duplicate records, assets that no longer exist and assets never registered. That discovery is a benefit of the process rather than a failure of it, but must be planned for and staffed.

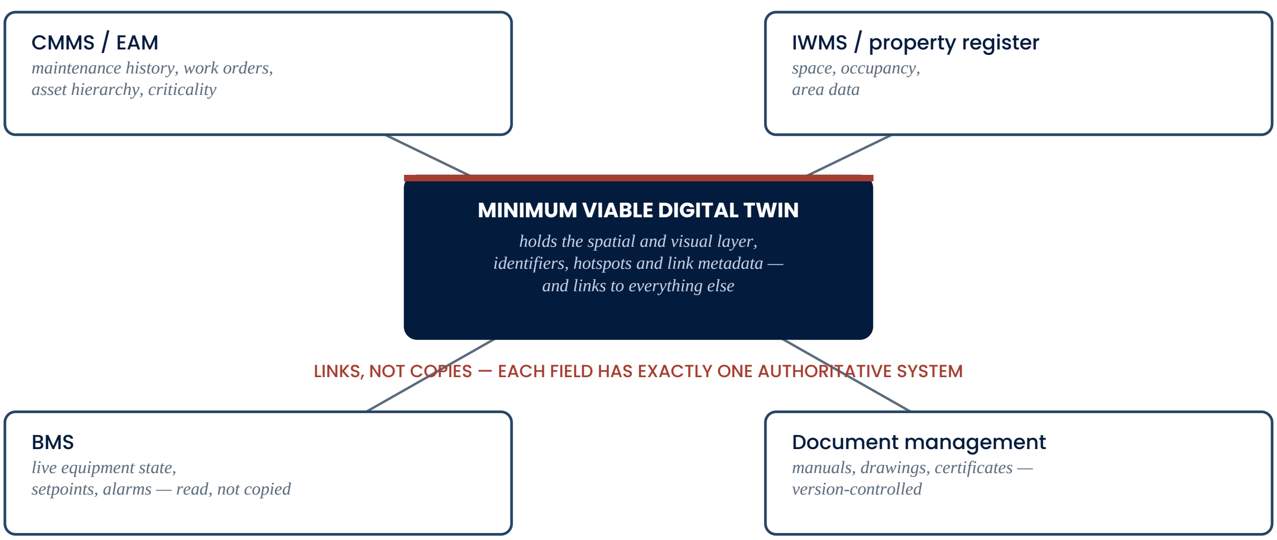

Figure 6. Reference architecture showing which system is authoritative for each type of information, and where the twin links rather than stores.

8. Pilot roadmap

A pilot exists to produce evidence, not coverage. Scope should be bounded to a single building, floor, system or asset category and to no more than three use cases — a constraint that keeps the effort finishable and the result interpretable.

PREPARE
steps 1–3

1 Select use cases
2 Establish scope
3 Create the location hierarchy

CAPTURE & CONNECT
steps 4–11

4 Capture spatial information
5 Process and quality-check
6 Identify in-scope assets
7 Create / validate identifiers
8 Link assets to records
9 Link documents
10 Establish access controls
11 Establish synchronization rules

OPERATE & PROVE
year one onward

12 Train users, by task
13 Measure adoption + outcomes
14 Expand only after evaluation

QUALITY GATE after step 5 — recapture before publishing

EVALUATION GATE after step 13 — controls expansion

Figure 7. The fourteen-step implementation sequence, showing the two decision gates that control expansion.

Suitable pilot use cases share three characteristics: they recur often enough to measure within a quarter, they involve a cost already visible in operational records, and they can be improved by information alone rather than by additional staff or capital.

- **Remote work-order triage** — determining from the desk what a reported fault involves, which asset is affected and what the responding technician should bring.
- **Equipment identification** — confirming which unit is which, particularly where labelling is inconsistent or several similar units share a room.
- **Vendor quotation support** — providing access routes, clearances and equipment context so quotations can be prepared without a site visit.
- **Technician onboarding** — orienting new or covering staff to layout, plant rooms and isolation points before their first shift.
- **Emergency planning** — documenting isolation points, shut-off locations, access routes and egress for planning and drills.
- **Asset inventory verification** — confirming that the asset register matches the physical building, and resolving the differences.
- **Condition documentation** — establishing a dated visual baseline for condition assessment, dilapidation, insurance or capital planning.

8.1 Establishing baselines before implementation

A baseline measured after implementation is an estimate, and estimates cannot support a business case. Baselines should be established over a defined period — a quarter usually suffices for high-volume activities — using data the organization already produces.

Work-order records support several measures directly: the count and proportion requiring a preliminary identification visit, return visits attributable to incorrect parts or asset identification, and elapsed time from work-order creation to first productive intervention. Search time can be measured through a short, time-boxed sampling exercise in which a few technicians record how long they spend locating asset information; published research has used self-reported estimates, which tend to overstate [3]. Vendor scoping visits can be counted from purchase orders or visit logs, and onboarding effort as supervised hours before a technician works independently.

Each baseline should record its method, period and source so the later comparison is genuinely like for like. Where a baseline cannot be measured, the business case should say so as an assumption rather than present it as a finding.

9. Building a defensible business case

Generic digital-twin savings percentages should not be used as evidence of value in a particular existing building. Published figures overwhelmingly describe sensor-dense or new-construction deployments, are frequently vendor-sponsored and rarely disclose their method. A business case built on borrowed percentages fails at the first informed question; one built on local measurement, transparent assumptions and stated uncertainty survives it.

9.1 The formulas

> **Annual operational benefit** = avoided site visits + reduced search time + reduced repeat work + faster vendor scoping + improved inspection productivity + reduced onboarding effort + reasonably quantifiable risk reduction
>
> **Net annual benefit** = annual operational benefit − annual operating and administration costs
>
> **Simple payback (years)** = initial implementation cost ÷ net annual benefit

Each term resolves to a quantity multiplied by a unit value: visits avoided by the fully loaded cost of a visit including travel; hours saved by a fully loaded hourly rate; return visits avoided by their cost; inspections gained at the inspector rate; supervised hours avoided at the supervising rate. Faster vendor scoping usually appears as reduced quotation lead time and fewer chargeable mobilizations.

Risk reduction is the term most often abused. Quantify it only where an insurer, a regulator or the organization's own risk register already attaches a value — a documented premium adjustment, a published penalty schedule, the recorded cost of a past incident better information would have prevented. Otherwise record it as a non-financial benefit and leave it out of the arithmetic.

Annual operating and administration costs must include the items pilots routinely omit: subscription or licence fees, recapture and update effort, the data owner's time, storage, training for new staff, and administration of access control.

9.2 Separating measurements, assumptions and estimates

Every line should carry one of three labels, visible to whoever approves it. A **measurement** is drawn from a system of record over a stated period. An **assumption** is a value chosen because it cannot be measured, stated with its basis and sensitivity. An **estimate** is a projection from a measurement, stated with its method. A case in which most benefit lines are assumptions is not necessarily wrong, but it is a case for a pilot rather than a portfolio rollout, and should be presented as such.

> *"A case in which most benefit lines are assumptions is a case for a pilot, not a portfolio rollout."*

Non-financial benefits should be listed separately and not converted into currency: knowledge retained when experienced staff leave, better evidence for insurance and dispute resolution, improved emergency preparedness, reduced dependence on individual memory, and greater capability to specify information requirements for future capital projects.

9.3 Illustrative scenario

> **This scenario is illustrative and hypothetical.** It uses assumed values to demonstrate the arithmetic. It does not describe a completed project by any organization, and its figures must not be cited as evidence of outcomes.

Assume a 12,000 ft² office and laboratory building with 180 in-scope mechanical and electrical assets, operated by four technicians and one manager. Assumed values: a fully loaded technician rate of 60 currency units per hour; an identification-only visit consuming 1.5 hours; a baseline of 240 such visits per year, measured from work-order records over one quarter and annualized; and a 40 percent reduction after implementation.

Avoided site visits: 240 × 0.40 × 1.5 hours × 60 = 8,640 per year (one measurement, one assumption). Reduced search time: 20 minutes saved per technician per working day across four technicians over 220 days = 17,600 per year (assumption; the most sensitive line, and one to replace with measurement before any expansion decision). Reduced vendor mobilizations: 12 avoided chargeable visits at 250 each = 3,000 per year (assumption). Illustrative annual operational benefit: 29,240.

An assumed annual operating cost of 9,000 — subscription, recapture of changed spaces, the data owner's time — gives an illustrative net annual benefit of 20,240. Against an assumed initial cost of 34,000 covering capture, identifier reconciliation, linking and configuration, simple payback is about 1.7 years. The purpose is the structure rather than the result: two of three benefit lines are assumptions, so the honest conclusion is that the case justifies a measured pilot, not a portfolio commitment.

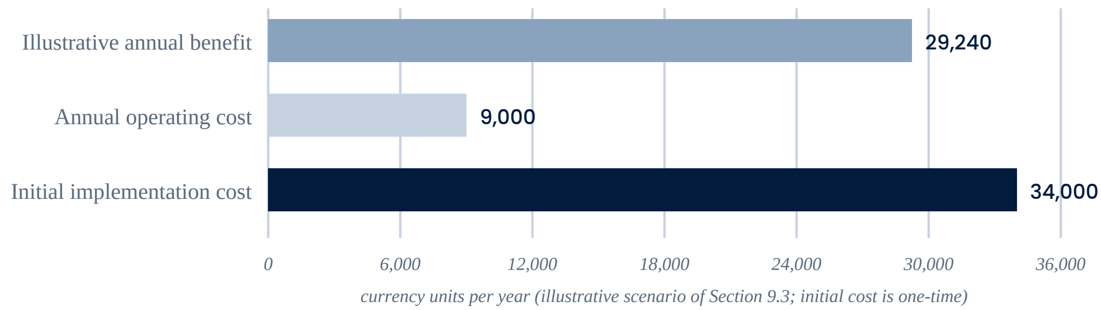

Figure 8. The illustrative scenario of Section 9.3 as arithmetic: the annual benefit the assumptions produce, against annual operating and one-time implementation costs, with simple payback of about 1.7 years. Every input is a stated assumption or a single measurement — the structure, not the result, is the point.

10. Governance, privacy and cybersecurity

A navigable visual record is a security-relevant asset. It documents door hardware, camera coverage, panel locations, communications rooms, isolation points and access routes, in a form needing no technical skill to interpret. Guidance for built-environment digital twins treats trustworthiness as a founding requirement alongside purpose and function [10], and operational-phase information management includes a security-minded approach as a distinct discipline [11]. These controls should be settled before capture, not retrofitted after publication.

10.1 Ownership, roles and access

A named individual should own the twin as an information asset, distinct from whoever administers the platform. Roles should be defined by task — view, annotate, edit links, manage users, export — and granted on least privilege. Contractor access should be time-bounded, scoped to relevant areas and reviewed on a schedule rather than on departure alone. Restricted rooms and security infrastructure should be registered with the security function before capture and either excluded or separately permissioned.

10.2 Privacy

Capture in occupied buildings records people, personal workspaces, visible screens and tenant property. A written privacy rule should state what is captured, what is redacted, who may view it and how long it is retained, and should be communicated to occupants and tenants before capture rather than discovered afterwards. Redaction should be verified as a quality-control step, and unredacted originals governed under the same retention rules. Where healthcare, education, custodial or regulated environments are involved, or personal data protection law applies, the rule should be reviewed by the function accountable for data protection in that jurisdiction.

10.3 Integrity, continuity and exit

Version control and audit history should record what changed, when and by whom, for imagery and asset links alike. Retention should state how long superseded imagery is kept — often longer than expected once it serves as condition or dispute evidence. Backups should be tested by restoring, not merely scheduled, and continuity planning should establish how information is reached when the platform is unavailable, including whether an offline view is needed for emergency response.

Exit provisions deserve as much attention as implementation. The organization should hold in writing: ownership of the captured imagery and derived data; the right to export imagery, metadata, identifiers and links in a documented, non-proprietary format at any time and on termination; the hosting jurisdiction and any restrictions arising from it; and the provider's obligations on termination, including data return and deletion. Authentication should align to the organization's identity standard, with single sign-on where available, and mobile access should be tested on the devices and networks technicians actually use.

Governance should also reach into future projects. Every renovation, replacement and fit-out is a chance to update the twin at marginal cost, and a chance missed if closeout requirements omit it. Adding capture and asset-data update to standard closeout documentation is the single most effective measure against ageing.

> *"Adding capture and asset-data update to standard closeout documentation is the single most effective measure against ageing."*

11. Risks and failure modes

Programmes of this kind rarely fail suddenly. They degrade, usually through one of the mechanisms below, each of which has an early indicator a manager can watch for.

Risk	Early indicator	Control
Describing a virtual tour as a digital twin	Stakeholders expect current information from a record that has no synchronization commitment.	Use the maturity stages explicitly in internal communication; state the synchronization frequency wherever the representation is presented.
Capturing everything without a use case	Capture scope is defined by building area rather than by operational need.	Require an outcome statement with a baseline before any capture is authorized.
Over-modelling	Effort is spent on geometry that no defined decision requires.	Test each modelling request against a named decision and the cost of that decision being wrong.
Under-documenting critical assets	Occupied areas are well covered; plant rooms, roofs and risers are thin.	Weight capture density by criticality, not by accessibility or appearance.
Stale information	No recapture since the initial programme; project closeouts do not reference the twin.	Publish the synchronization schedule; embed capture tasks in closeout procedures; display the last-verified date to users.

Risk	Early indicator	Control
Duplicate data	Asset attributes are edited in the twin platform.	Publish the information-ownership map; disable editing of fields mastered elsewhere.
Incorrect identifiers	Hotspots resolve to the wrong record; technicians report mismatches.	Reconciliation reporting with a match-rate target and a managed exception list.
No data owner	Corrections are reported but not actioned; nobody can approve a change.	Name an owner per building or portfolio segment in the governance register.
Poor technician adoption	Usage concentrates in office roles; field usage is negligible.	Task-based training; mobile and offline testing on real devices; include field staff in scope decisions.
Vendor lock-in and proprietary formats	Export produces a viewer package rather than usable data.	Test a full export during procurement, before contract signature, and confirm a second party could work with it.
Unverified accuracy	Dimensions are taken from panoramic imagery for procurement decisions.	State measurement limitations in the interface; specify accuracy formally where measurement matters [13].
Privacy and cybersecurity exposure	Broad pilot access is never narrowed; restricted areas appear in general collections.	Access matrix reviewed on a schedule; restricted-area register agreed before capture; periodic access audit.
Unsupported return claims	The business case cites published percentages rather than local measurement.	Label every line as measurement, assumption or estimate; report sensitivity on the largest assumptions.
Future changes not incorporated	A renovation completes and the twin still shows the previous configuration.	Make capture and data update a signed closeout deliverable on every project above a defined threshold.

Table 5. Risks, early indicators and controls.

Three deserve emphasis for small teams. A programme with no named data owner will decay regardless of how well it was built. One whose export has never been tested is a single procurement decision away from being unrecoverable. And one field technicians do not use has not failed technically — it answered a question nobody was asking, which no platform capability will correct.

12. Procurement questions

Platform selection is an information-governance decision rather than a software purchase, because the durable risks concern ownership, portability and continuity rather than features. The full question set appears in Appendix E, written to be pasted into a request for proposal (RFP).

Four questions carry disproportionate weight. **Who owns the data, and what exactly is exported?** A demonstrated export during evaluation — not a contractual promise — is the only reliable answer, and the test is whether a different provider could work with the exported material. **What is authoritative, and how is integration achieved?** Documented application programming interfaces and support for open exchange schemas indicate a platform designed to connect rather than to contain [9]. **Who can see what?** Role-based access, single sign-on, audit trails and controls for sensitive areas should be demonstrated, not described. **What happens on termination?** Data return, deletion, format and timescale should be specified in the contract rather than negotiated under pressure.

Accuracy claims require particular scrutiny. Any statement about measurement capability should be tied to a recognized accuracy framework with a stated confidence level, and the supplier should say plainly which deliverables are measurement-grade and which are not [13]. A supplier unwilling to state the limitations of its own output is describing a marketing position rather than a technical one.

13. Conclusion

Existing buildings do not need a complete model reconstruction or a dense sensor network to begin. They need a defined operational outcome, navigable spatial context, reliable location and asset identities, connections to authoritative records and workflows, a documented synchronization and change-management process, and governance that makes the result trustworthy. Those six requirements are achievable within a normal operating budget for a bounded scope, and they separate a digital twin from a record of how a building once looked.

The distinction that most often goes unmade is worth restating. A 360° virtual tour is a reality-capture layer with real and immediate operational value. It is not a digital twin. It becomes part of one when identity, connection and synchronization are added, and describing it accurately before that point protects the credibility of the programme that follows.

The recommendation is deliberately narrow. Start with one operational problem, one governed spatial baseline and one trustworthy set of linked asset information. Measure the outcome against a baseline established before implementation. Add geometry, sensors, analytics and automation only when they support a demonstrated operational need. The organizations that reach analytical and predictive capability in existing buildings will generally be those that first made a small amount of information reliable, current and governed.

For a manager persuaded of the argument, the first month requires no procurement at all. Choose the one recurring problem the pilot will address and write its outcome statement. Pull the baseline from records the organization already produces — work orders, purchase orders, onboarding logs. Walk the identifier hierarchy from portfolio to space against door signage and the maintenance system, and record every conflict found. Register the restricted areas with the security function. None of this commits a budget, and all of it must be done before capture regardless of which platform, method or vendor is eventually selected.

The second month is the pilot described in Section 8, bounded to one building and at most three use cases, with the scorecard in Appendix D completed as the work proceeds rather than reconstructed at the end.

The work rewards the order in which it is done. Identity before capture, connection before expansion, measurement before claims — each step makes the next one cheaper, and each one skipped is paid for later at a higher price. That is the practical meaning of *minimum viable*: not the least that can be bought, but the least that can be trusted.

> *“Identity before capture, connection before expansion, measurement before claims — not the least that can be bought, but the least that can be trusted.”*

The test to apply before any expenditure. Name the operational problem, the measure and the baseline. State the synchronization frequency and fidelity the decision actually requires. Identify the authoritative system for every field the twin will show. If any of the three cannot be stated, the programme is not ready for capture — and no platform purchase will change that.

Appendix A — Minimum viable digital twin checklist

A representation qualifies as a minimum viable digital twin when every item below can be evidenced. Any unchecked item should be recorded with an owner and a date.

Requirement	Checklist item	Evidence to produce
Defined operational outcome	A written outcome statement exists for each use case (maximum three).	Signed one-page statement per use case.
	A baseline was measured before implementation, not estimated afterwards.	Baseline data with source, method and period.
	An accountable owner is named for each outcome.	Governance register entry.
Navigable spatial context	All in-scope spaces are captured and navigable by an unfamiliar user.	Capture log with coverage list.
	Capture positions are indexed to space identifiers.	Position-to-space index.
	Inaccessible spaces are recorded with a reason and an owner.	Recapture list.
Unique identities	A published identifier scheme covers portfolio through asset.	Identifier standard document.
	Identifiers are opaque, persistent and never reused.	Written rule plus decommissioning procedure.
	In-scope assets are physically labelled and field-verifiable.	Labelling completion report.
	Match rate between tour, records and field labels is measured.	Reconciliation report and exception list.
Connections	Each field is assigned to exactly one authoritative system.	Information-ownership map, dated.
	Assets resolve from the twin to their authoritative record.	Traced sample of at least twenty assets.
	At least one workflow is reachable from the twin.	Demonstration record.

Requirement	Checklist item	Evidence to produce
Synchronization	A synchronization schedule states frequency, fidelity, trigger and owner per element.	One-page schedule, dated.
	Project closeout includes capture and data update as a signed deliverable.	Closeout checklist template.
	Updates have occurred on schedule since go-live.	Capture log and update history.
Governance	A named data owner is appointed.	Governance register.
	Role-based access is defined and reviewed on a schedule.	Access matrix with review dates.
	Restricted areas were agreed with the security function before capture.	Restricted-area register.
	Privacy rule is written, communicated and applied.	Privacy rule plus redaction quality-control record.
	A full data export has been performed and verified as usable.	Export test record.
	A backup restore has been tested.	Restore test record.

Appendix B — 360° capture field checklist

Before mobilizing

- Capture scope agreed in writing, including explicit exclusions.
- Building, floor and space identifiers established and reconciled with signage and the maintenance system.
- Restricted-area register agreed with the security function.
- Privacy rule agreed and occupants or tenants notified.
- Access arranged for plant rooms, risers, roofs and locked spaces, with escort where required.
- Site safety requirements confirmed: permits, personal protective equipment, roof access, energized-equipment restrictions.
- Capture log, recapture list and naming convention prepared.

During capture

- Camera height consistent and recorded.
- Scene spacing close enough that the next position is visible from the current one.
- A capture position on each side of every doorway.
- Positions at every stair landing; lobby positions at each elevator floor.

- Multiple positions in each mechanical, electrical and telecommunications room.
- Rooftop positions showing access, curbs, drainage and unit orientation.
- Selective above-ceiling capture at known access points, where safe.
- Bracketed exposure in high-contrast spaces.
- Mirrors, glazing and reflective surfaces noted in the log.
- Capture outside occupied hours where practical; moving people and vehicles noted where not.
- Nameplate and serial-number photography for every in-scope asset, tied to the asset identifier.
- Detail photography of occluded equipment, valve tags and panel schedules.
- Every position logged with space identifier, date and operator.
- Every missed or inaccessible space logged on the recapture list with a reason.

Quality control before publication

- Stitching artefacts reviewed, particularly at doorways and near objects.
- Nadir and zenith coverage confirmed on every scene.
- Operator, tripod and equipment removed.
- Exposure and white balance consistent across each floor.
- Privacy redaction applied and verified.
- Restricted areas excluded or moved to a permissioned collection.
- Space identifiers verified against the hierarchy.
- Capture date visible to users on every scene.
- Superseded imagery retained and versioned by date.
- Recapture list assigned with owners and due dates.

Using this checklist

The checklist is sequenced deliberately. Every item under “before mobilizing” is cheaper to resolve at a desk than in the field: a missing identifier hierarchy discovered on site suspends the capture; discovered a week earlier, it is an email. The quality-control items are pass–fail gates, and publishing imagery that fails them costs more than the recapture it avoids, because the first impression the record makes on technicians decides whether it is used.

Where capture is contracted rather than performed in-house, attach the checklist to the scope of work and make the capture log, recapture list and quality-control confirmation named deliverables. A contractor asked for them at award will produce them; one asked at handover usually cannot.

Adapt the list to site conditions rather than abbreviating it, and file the completed checklist with the capture log. Together with the recapture list, they are the audit trail the spatial-context requirement in Section 3.2 calls for, and the evidence an incoming manager will look for first.

Appendix C — Minimum asset data dictionary

Field	Type	Required	Definition and rule	Example
portfolio_id	Text	Multi-site only	Identifier for the ownership or management grouping. Stable; never reused.	PF-NORTH
site_id	Text	Yes	Identifier for the physical site or campus.	ST-0042
building_id	Text	Yes	Identifier for the building. Must reconcile with signage and lease records.	BL-0042-C
floor_id	Text	Yes	Identifier for the level, following wayfinding convention.	FL-03
space_id	Text	Yes	Identifier for the room or space. Reconcile with door signage before capture.	RM-03-118
asset_id	Text	Yes	Portfolio-unique, opaque and persistent. Survives replacement of the physical unit.	AST-000184
asset_tag	Text	Yes	Human-readable field label applied to the physical asset.	AHU-3-02
asset_category	Controlled list	Yes	Classification for reporting. Use the organization's existing list.	Air handling unit
system	Controlled list	Yes	System the asset serves. Supports shutdown impact assessment.	HVAC — supply air
manufacturer	Text	Yes	As printed on the nameplate.	Recorded from nameplate
model	Text	Yes	As printed, including suffixes.	Recorded from nameplate
serial_number	Text	Yes	Identifies the physical unit. Changes on replacement.	Recorded from nameplate
criticality	Controlled list	Yes	Organization's existing criticality scale. Do not create a parallel scale.	1 (highest) to 4
maintainable	Boolean	Yes	Whether the asset generates planned maintenance.	TRUE
condition_status	Controlled list	Yes	Observed condition at last verification. Always paired with verified_date.	Fair
install_date	Date	Where known	Installation or replacement date.	2014-06 (estimated)
warranty_expiry	Date	Where applicable	Expiry date; record the party holding the obligation separately.	2027-03-31
warranty_holder	Text	Where applicable	Party responsible for warranty performance.	Installing contractor

Field	Type	Required	Definition and rule	Example
info_source	Controlled list	Yes	Origin of the record values.	Nameplate / drawing / submittal / technician report / estimate
data_confidence	Controlled list	Yes	Three levels only: verified, reported, estimated.	Verified
verified_date	Date	Yes	Date the record was last confirmed against the physical asset.	2026-04-18
data_owner	Text	Yes	Accountable role, not a department or individual name.	Facility Information Manager
linked_documents	Reference list	Yes	References to manuals, drawings and certificates. Store links, never copies.	DMS reference identifiers
linked_service_history	Reference list	Yes	Reference to the maintenance record for the asset.	CMMS asset record identifier

Field names are supplied as a working convention. Where the organization already operates a data standard, adopt that standard and map to it rather than introducing a second vocabulary.

Population guidance

Populate the dictionary in three passes rather than asset by asset. **Identity fields first** — portfolio through space come from the hierarchy agreed before capture, and asset_id and asset_tag from the reconciliation walk. **Nameplate fields second** — manufacturer, model and serial_number from the photographs taken during capture, transcribed exactly as printed. **Status and lifecycle fields last**, because condition, install dates and warranty terms take longer to trace and should not delay the first two passes.

Three rules keep the result trustworthy. data_confidence is never left blank — “estimated” is an acceptable value and an honest one, while an empty field cannot be distinguished from an unexamined one. verified_date changes only when someone has stood in front of the asset or its nameplate photograph, not when a record is edited. And controlled lists stay short: a category or criticality list that grows past what a technician can hold in mind stops being applied consistently, which defeats its purpose.

Where values are inherited from an existing register, carry them across with data_confidence set to “reported” until field verification upgrades them. Bulk-importing inherited values as “verified” is the fastest way to make the most useful field in the dictionary meaningless.

Fields may be added to the dictionary, but each addition should name its authoritative system before it is collected. The dictionary grows by the rule in Section 6.2 — one master per field — not by enthusiasm, and a field nobody is accountable for keeping current is better left out than left stale.

Appendix D — Pilot measurement scorecard

Complete the baseline column before implementation. Record the measurement method and source for every row; rows without a stated method should be labelled as assumptions in the business case.

Measure	Method and source	Baseline period	Baseline value	Post value	Category
Identification-only site visits per period	Work-order records, filtered by reason code				Measurement
Return visits attributable to wrong part or wrong asset	Work-order records, linked visits				Measurement
Elapsed time from work-order creation to first productive intervention	Work-order timestamps				Measurement
Technician time spent locating asset information	Time-boxed sampling by named technicians				Estimate
Vendor mobilizations for scoping	Purchase orders or visit logs				Measurement
Quotation lead time	Procurement records				Measurement
Inspections completed per period	Inspection records				Measurement
Supervised hours before independent working	Onboarding records				Measurement
Asset register match rate against field verification	Reconciliation report				Measurement
Twin usage by role	Platform usage reporting				Measurement
Proportion of in-scope assets with verified_date within 12 months	Data-quality report				Measurement
Open items on the recapture and exception lists	Capture and reconciliation records				Measurement

Evaluation gate. Expansion should be authorized only where at least one outcome measure has improved against its baseline, adoption among the intended field roles is demonstrated, the synchronization schedule has been met, and a full data export has been tested. Where an outcome measure has not improved, the more likely explanation is an incorrect use case than an incorrect platform.

Appendix E — Vendor and platform RFP questions

Data ownership and export

1. Who owns the captured imagery, derived data, identifiers, links and metadata during the term and after termination? Quote the contractual clause.
2. Provide a complete export from a representative account during evaluation. Identify every file format and state which are open and documented.
3. Confirm that exported material includes imagery, capture metadata, positions, identifiers, hotspot mappings, links and audit history.
4. Could a different provider ingest the exported material and reconstitute an equivalent service? Describe how.
5. Is export available on demand and without additional charge? State any limits on frequency or volume.

Integration and standards

6. Provide public documentation for all application programming interfaces, including authentication, rate limits and versioning policy.
7. Which open exchange schemas for facility and asset information are supported, in which direction, and to what extent?
8. Describe existing integrations with maintenance and workplace management systems, and state which were delivered as configuration rather than custom development.
9. How are identifiers mapped and maintained between systems, and what happens when a mapping breaks?
10. What is the notice period for breaking changes to interfaces?

Access, identity and audit

11. Describe the role model and demonstrate role-based access using test accounts.
12. Is single sign-on supported against the organization’s identity provider, and at what licence tier?
13. How are sensitive or restricted areas controlled, and can they be excluded from general access entirely?
14. What audit trail is retained, for which actions, for how long, and how is it exported?
15. How is external contractor access time-bounded and reviewed?

Versioning, accuracy and limitations

16. How is superseded imagery retained and version-compared over time?
17. State plainly which deliverables are measurement-grade and which are not.
18. Where measurement is offered, against which recognized accuracy framework is it specified, at what confidence level, and how is it validated?
19. What warnings are displayed to users where measurement is unreliable?

Hosting, continuity and termination

20. In which jurisdictions is data hosted and processed, and can hosting be restricted to a specified region?
21. Describe backup frequency, retention, and the last tested restore.
22. State the availability commitment and the remedy for failing it.
23. On termination, what is returned, in what format, within what period, and how is deletion evidenced?
24. If the provider ceases trading or is acquired, what continuity provisions apply?

Usability, accessibility and field conditions

25. Demonstrate the mobile experience on the device models and operating system versions the organization actually deploys.
26. Is offline or cached access available for areas without connectivity, such as basements and plant rooms?
27. State conformance with the organization's accessibility standard and provide a current accessibility conformance report.
28. What training is provided, and is it task-based rather than feature-based?

Commercial and long-term maintainability

29. State all recurring costs, including storage growth, additional users, additional buildings and recapture.
30. What is the price protection at renewal?
31. What is required to maintain the information if the subscription lapses?
32. Provide two reference customers operating existing buildings of comparable type who have completed at least one recapture cycle.

References

Bracketed numbers in the text correspond to the entries below.

[1] Digital Twin Integration Subcommittee. (2024, June 25). *Digital Twins for the Built Environment: A Position Paper on Integrating BIM and Digital Twin.* National Institute of Building Sciences, Digital Technology Council.

[2] IFMA Information Technology Community & Autodesk. (2025, March). *Advancing Facilities Management with a Digital Thread.* IFMA IT Community. (Citing *Optimizing Building Management with a Lifecycle Approach*, 2023 survey.)

[3] IFMA Information Technology Community & Autodesk. (2025, March). *Advancing Facilities Management with a Digital Thread.* IFMA IT Community. (Citing Autodesk & FMI, *Harnessing the Data Advantage in Construction*, and 2024 conference polls.)

[4] Gallaher, M. P., O'Connor, A. C., Dettbarn, J. L., & Gilday, L. T. (2004). *Cost Analysis of Inadequate Interoperability in the U.S. Capital Facilities Industry* (NIST GCR 04-867). National Institute of Standards and Technology.

[5] Digital Twin Integration Subcommittee. (2024, June 25). *Digital Twins for the Built Environment*, Section 3, Positions on Use Cases. National Institute of Building Sciences.

[6] Digital Twin Consortium. (2024, October 3). *Definition of a digital twin.* Digital Twin Consortium.

[7] Ghorbani, Z., & Messner, J. I. (2024). A categorical approach for defining digital twins in the AECO industry. *Journal of Information Technology in Construction (ITcon)*, 29(10), 198–218. As adopted in the NIBS position paper.

[8] International Organization for Standardization. (2020). *ISO 19650-3:2020, Organization and digitization of information about buildings and civil engineering works, including building information modelling — Information management using building information modelling — Part 3: Operational phase of the assets.*

[9] buildingSMART International. *Industry Foundation Classes (IFC)* (ISO 16739-1); and National Institute of Building Sciences, *Construction to Operations Building information exchange (COBie)*, National BIM Standard – United States.

[10] Bolton, A., Butler, L., Dabson, I., Enzer, M., Evans, M., Fenemore, T., et al. (2018, December). *The Gemini Principles: Guiding values for the national digital twin and information management framework.* Centre for Digital Built Britain.

[11] International Organization for Standardization. (2020). *ISO 19650-5:2020, Organization and digitization of information about buildings and civil engineering works, including building information modelling — Information management using building information modelling — Part 5: Security-minded approach to information management.*

[12] Digital Twin Integration Subcommittee. (2024, June 25). *Digital Twins for the Built Environment*, Section 1.1, Categories of a Digital Twin, and Section 3.4.3, Physical–Digital Relationship. National Institute of Building Sciences.

[13] U.S. Institute of Building Documentation. (2025, January). *Level of Accuracy (LOA) Specification*, Version 3.1.

Further reading

International Organization for Standardization & International Electrotechnical Commission. (2023, November). *ISO/IEC 30173:2023, Digital twin — Concepts and terminology.*

Digital Twin Consortium. (2024, October 1). *Definition of digital thread.* Digital Twin Consortium.

National Academies of Sciences, Engineering, and Medicine. (2024). *Foundational Research Gaps and Future Directions for Digital Twins.* The National Academies Press.

International Organization for Standardization. (2014). *ISO 55000:2014, Asset management — Overview, principles and terminology.*

Abbreviations

Abbreviation	Meaning
BIM	Building information model / modelling
BMS	Building management system
CMMS	Computerized maintenance management system
COBie	Construction to Operations Building information exchange
EAM	Enterprise asset management
IFC	Industry Foundation Classes
IWMS	Integrated workplace management system
LiDAR	Light detection and ranging
LOA	Level of Accuracy (USIBD specification)
MEP	Mechanical, electrical and plumbing
NIBS	National Institute of Building Sciences
RFP	Request for proposal
SLAM	Simultaneous localization and mapping

A note on this edition

This is Version 1.0, the first published edition. Corrections and suggestions for future editions are welcome and may be directed to the publisher; substantive changes in any later edition will be recorded in the document history below. The standards cited in the References are available from their publishing bodies, and the position papers and guidance documents were publicly available at the time of publication.

Document history

Version	Year	Change
1.0	2026	First publication.

About the author

Nikolai Hanov has more than 20 years of facilities and operations experience across multi-site portfolios, covering maintenance, capital planning, vendor management, facility data and operational improvement. His work focuses on making facility information usable by the people who operate buildings, rather than on technology adoption for its own sake. He is the author of *Strategic Facility Management: How Buildings Become Profit Centers, Not Cost Centers.* He is affiliated with Levaru.

About Levaru

Levaru is active in the built-environment information field. Its work concerns how information about buildings is captured, structured and kept usable by the people who operate them. This paper is published as part of that work and, as stated in the vendor-neutrality disclosure on page 2, names no commercial platform and describes capability categories rather than implementations.

LEVARU

Research on the operational record of buildings

Also in this series

The Visual Asset Register · January 2026

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