

RESEARCH REPORT

The Ontology Imperative and the Awareness Threshold

What the evidence supports, and what a non-technical commercial real estate workforce must be told first.

Building Lifecycle Management Initiative

Executive Summary.....	4
The Question, and Why Prior Work Does Not Answer It.....	6
Part I: Is There an Imperative?	6
2.1 The government and standards record.....	6
2.2 The 15.8 billion dollar figure, stated correctly	7
2.3 What deployment actually costs	7
2.4 The portability evidence, which cuts both ways.....	8
2.5 The grounding evidence, and its limits	8
2.6 The buildings ceiling, which is the finding that matters.....	9
2.7 When the system acts.....	9
Part II: Why This Is a Facilities Problem.....	10
3.1 One word, several meanings, five systems.....	10
3.2 Handover, where the record is strongest	11
3.3 Time, which is the cheapest fix and the most expensive omission.....	12
Part III: The Audience.....	12
4.1 What is known, and what is asserted.....	12
4.2 The counter-signal that deserves airtime.....	13
4.3 False familiarity, which is the actual obstacle	13
Part IV: The Awareness Tier	14
5.1 Where the scale comes from	14
5.2 A boundary test	15
5.3 What must be communicated at Awareness	15
5.4 What must be handled, not merely avoided	16
5.5 What does not belong at Awareness.....	17
Part V: The Honest Caution.....	17
Conclusion.....	18
Gaps in the Evidence Base	19
Works Cited.....	20
About BLMI.....	25

Disclaimer 25

AI Use Transparency 25

Executive Summary

The argument that formal semantics has become a precondition for artificial intelligence in buildings, rather than a quality improvement, is well supported by primary government sources and peer-reviewed literature. It is also more qualified than its advocates generally concede, and the qualification changes what should be said to a non-technical audience.

Six findings carry this report.

First, the precondition claim holds, and it holds on government authority rather than vendor assertion.

The National Institute of Standards and Technology states that harvesting building data "remains labor-intensive, requiring manual mapping of data sources to the needs of each application," and that this "hinders scalability, increases costs, and delays deployment." The US Department of Energy states that building software "is not plug and play." Neither is selling anything.

Second, formal semantics is necessary but demonstrably not sufficient. The Mortar testbed applied unmodified analytics applications across 90 semantically modeled buildings. The best-performing application ran on 67 percent of them. Most ran on fewer than 40 percent. A companion study found that modeling variance within a single shared schema is enough to break queries outright. Anyone claiming that adopting a semantic standard makes analytics portable is claiming more than the evidence supports.

Third, the buildings-specific result for artificial intelligence is uncomfortable and should be stated plainly. The only rigorous published benchmark of large language models answering questions over building knowledge graphs found that even the best-performing model, on a combined measure of how often its extracted answers were both correct and complete, reached only 0.38 out of a possible 1.0, with exact-match accuracy near zero, and confirmed that models invent classes and relationships that do not exist in the ontology. That failure was worse against the newer ASHRAE semantic standard than against the older one. The correct conclusion is not that semantics does not matter. It is that semantics is the floor, the residual error rate is far too high to license autonomous action today, and the argument "add the graph and the agent can act" is contradicted by the best available evidence.

Fourth, the risk profile changes materially when a system acts rather than advises, and that shift now has regulatory expression. European law classifies artificial intelligence acting as a safety component in the supply of heating or electricity as high risk, justified explicitly by consequence of failure, and requires that a human overseer be able to override or halt the system. Industry consensus names excessive agency as a distinct vulnerability class. Peer-reviewed security research establishes that a system combining retrieval with the ability to act inherits its trust boundary from its data sources.

Fifth, the workforce premise is real but is not the premise commonly cited. There is no credible measurement of data literacy in the facilities management or commercial real estate workforce specifically. Every circulating figure originates in general-workforce, self-reported, vendor-sponsored surveys. What is well evidenced is different and more useful: workforce capability, not technology availability, is the named constraint on adoption. Only 33 percent of the commercial real estate workforce feel adequately trained on artificial intelligence, and only 5 percent of firms report achieving their stated objectives. The median United States facilities manager is 53.6 years old.

Sixth, and most consequential for the communication problem, the obstacle at the Awareness tier is not incomprehension. It is false familiarity. Cognitive research establishes that a superficially similar new concept reactivates a learner's existing and partially wrong understanding, and that unless the old belief is surfaced in the same moment as the correction, the learner encodes nothing while feeling that learning occurred. A practitioner who believes a naming convention is an ontology will not return for the next tier. The Awareness message must therefore name the misconception rather than assert past it.

The report closes with five items that belong at the Awareness tier, a boundary test for what does not, and the honest caution that awareness communication reliably moves knowledge and does not reliably move behavior.

The Question, and Why Prior Work Does Not Answer It

Two Building Lifecycle Management Initiative research products already map this terrain. A July 2026 scan examined eleven standards addressing location, space, and asset hierarchy and concluded that no formal standard defines a complete hierarchy from portfolio to component, that the Open Standards Consortium for Real Estate functions as an integration layer rather than a competitor, and that the gap is translation rather than invention. An August 2026 landscape assessment mapped which semantic, taxonomy, and ontology efforts exist across the building lifecycle, where they overlap, and where they are heading.

Both answer which standards exist. Neither answers whether the industry is obliged to adopt one, and neither addresses how to explain that obligation to people who do not work with data.

The commissioning input for this project was a synthesis of eleven ontology sources published between January and September 2026. That corpus is argumentative rather than empirical. Ten of its eleven sources are essays. The eleventh reports a build whose own author declines to call it an ontology. Three of its authors hold direct commercial positions in the conclusions they argue. The corpus states the imperative persuasively and evidences it thinly.

This report supplies the evidence, and reports where the evidence declines to cooperate.

Part I: Is There an Imperative?

2.1 The government and standards record

The strongest authority for the precondition claim comes from bodies with nothing to sell.

The National Institute of Standards and Technology runs a program called Building Digitization and Semantic Interoperability. Its published problem statement is that building data "come from diverse sources and span all phases of the building lifecycle," and that harvesting it "remains labor-intensive, requiring manual mapping of data sources to the needs of each application. This process hinders scalability, increases costs, and delays deployment." The program's stated objective is standardized machine-readable semantic models "to advance AI-enabled analytics, automation, and control." The institute positions semantics as the enabler of artificial intelligence, not as an optimization of it.

The US Department of Energy states the same barrier in operational terms. Buildings and their systems "are bespoke," so control and analytics software "must be configured individually for each building." Building software "is not plug and play, that is, self-configuring and installing, in the way that consumer software is." The department has committed five million dollars to date across national laboratories and universities to close that gap, with a further 1.3 million dollars directed specifically at the ASHRAE semantic data model work.

That standard, ASHRAE Standard 223P, exists "to define formal knowledge concepts and a methodology to apply them to create interoperable, machine-readable semantic frameworks for representing building automation and control data." It completed advisory public review in March 2024 and first public review in August 2025. As of early September 2026 no published edition had been located. Its behavioral counterpart, ANSI/ASHRAE Standard 231-2026, which formalizes control sequences rather than system structure, reached publication first. That sequencing is worth noting: the industry has settled how to state unambiguously what a system should do before it has settled how to state unambiguously what the system is.

Neither NIST nor the Department of Energy quantifies the cost of the gap. That absence matters, because the figure most often cited to fill it does not measure what people think it measures.

2.2 The 15.8 billion dollar figure, stated correctly

In 2004 NIST published a cost analysis finding 15.8 billion dollars in annual interoperability costs for the United States capital facilities industry. The figure is real, carefully derived, and self-described as conservative. It is also persistently misused.

Three qualifications govern its use. It is expressed in 2002 dollars and has not been updated in 22 years. It covers interoperability failure across design, construction, fabrication, and operations for the whole capital facilities industry, and it does not isolate building automation metadata, terminology conflict, or any other specific cause. Tenants were excluded from the study entirely.

What the figure does establish, and what is genuinely useful, is where the cost lands. Of the total, 9.09 billion dollars, or 57 percent, was attributed to the operations and maintenance phase, and roughly two-thirds of the whole was borne by owners and operators. The cost mechanism is mundane: files "had to be reentered manually into multiple systems," alongside "duplication of business functions" and searching paper archives. Whatever the number is today, its shape has held. The party that pays for poor definitions is the party that operates the building, and it pays over decades rather than at handover.

2.3 What deployment actually costs

Better-scoped economics come from Lawrence Berkeley National Laboratory. A 2022 study covering 113 organizations and more than 567 million square feet of commercial floor area found fault detection and diagnostics tools costing approximately nine dollars per point, or 13,000 dollars per building, with recurring costs of four dollars per point. A companion study across 550 buildings found median whole-building savings of eight percent and ongoing labor of eight dollars per point per year against a typical 1,300-point system.

Two things stand out. The savings are real and modest, at eight to nine percent median, which is far below the twenty to fifty percent commonly claimed in marketing material. And the studies name configuration duration, "several months," as a key barrier, alongside the labor of integrating diverse meter and building automation data.

Neither study breaks out how much of that cost is metadata work. That decomposition does not exist in the published record. An analyst firm asserts that up to 75 percent of engineering effort goes to making

existing systems legible to the analytics layer, but publishes no methodology, sample, or definition, and sells the report the claim appears in. This report records that assertion as an attributed industry claim and declines to treat it as a measurement.

2.4 The portability evidence, which cuts both ways

The clearest test of whether semantics delivers what its advocates promise is the Mortar testbed: 90 buildings, more than nine billion data points, each building modeled in a shared schema. Researchers then ran unmodified analytics applications across the whole set and counted how many buildings each one worked on.

The results are sobering. The best-performing application, a virtual coil meter, ran on 60 of 90 buildings, or 67 percent. A rogue zone temperature check reached 62 percent. Baseline calculation reached 37 percent, stuck damper detection 33 percent, and energy disaggregation 16 percent. Two applications ran on fewer than four buildings each.

A follow-on study explains part of the shortfall and is arguably more important. Across 50 building representations in the same schema, queries succeeded on some buildings and failed on others because of "differing modeling choices even when representing identical building systems." A shared vocabulary does not by itself produce shared models. Two competent modellers describing the same air handling unit can produce graphs that a single query cannot both read.

A separate survey of semantic modeling for building energy management found that existing ontologies "represent more maturely physical building entities than they represent abstract operational concepts," and reported that the necessity to extend an off-the-shelf ontology ran as high as 90 percent across reviewed use cases.

The honest reading is this. Semantic modeling is a necessary condition. Every serious platform effort builds on one, and both NIST and the Department of Energy identify manual mapping as the scaling barrier. But the published record does not support the claim that adopting a semantic standard has been shown to reduce deployment cost at portfolio scale, nor that it makes analytics portable. The most rigorous cross-building fault detection demonstration in the literature covers two museums.

2.5 The grounding evidence, and its limits

The case that artificial intelligence needs grounding is strong in principle and mixed in measurement.

In principle it is settled. Language model output "is not grounded in communicative intent, any model of the world, or any model of the reader's state of mind," and fluency causes readers to impute meaning that is not present. NIST enumerates confabulation, "production of confidently stated but erroneous or false content," among twelve generative artificial intelligence risks, and names data provenance as a control.

In measurement it is contested. The most cited enterprise result found that a large language model answering questions directly over a relational database scored 16.7 percent accuracy, while the same questions over a knowledge graph representation of the same database scored 54.2 percent. The figure that should carry more weight is buried in the detail: on the highest schema complexity, the model scored

zero without the ontology and roughly 36 to 39 percent with it. The study was authored by employees of a knowledge graph vendor, its baseline is deliberately naive zero-shot prompting rather than engineered query generation, and its domain is insurance.

Against that, an independent systematic evaluation found that graph-based retrieval does not universally beat conventional retrieval. On single-hop question answering, conventional retrieval won outright. The graph advantage appeared only on multi-hop reasoning, and it was bought at roughly fifty-seven times the index construction cost. A widely cited industry graph retrieval paper reports win rates of 72 to 83 percent, but those measure comprehensiveness and diversity as judged by another language model. Neither is correctness, and the result should not be quoted as an accuracy finding.

2.6 The buildings ceiling, which is the finding that matters

Everything above concerns other domains. The only rigorous benchmark of large language models operating over building knowledge graphs was published in November 2025. It tested five current models against 188 questions over four real buildings, three modeled in Brick and one in ASHRAE 223P.

Top-performing systems, on a combined measure of how often extracted answers were both correct and complete, reached a maximum of 0.38 out of a possible 1.0. Exact-match accuracy was near zero across every method tested. The study confirmed ontology hallucination empirically, "where the model invents relationships or classes that do not exist," and found the problem worse against the newer ASHRAE standard than against the older schema. It also found that methods excelling on general knowledge graphs "do not transfer effectively to the specialized vocabularies" of building operations.

A companion requirements paper from the same research community had already named the underlying condition: models "still struggle with the depth of knowledge required to accurately interpret building-specific data," and the field lacks "comprehensive evaluation datasets, standardized accuracy metrics, and benchmarks."

This is the result that should govern how the industry talks about the topic. It does not weaken the case for formal semantics. It strengthens the case while destroying a particular version of it. The version it destroys is the one that runs: build the ontology, connect the agent, let it act. Measured against the best available evidence, a system reading a well-modeled building graph gets roughly a third of the answers usefully right. That is a research frontier, not an operating capability, and it is nowhere near the accuracy required to license unsupervised action on physical plant.

2.7 When the system acts

The distinction between a system that advises and one that acts is not rhetorical. It now has legal and consensus expression.

European Regulation 2024/1689 classifies as high risk any artificial intelligence system "intended to be used as safety components in the management and operation of critical digital infrastructure, road traffic, or in the supply of water, gas, heating or electricity." The stated justification is consequence of failure: such failure "may put at risk the life and health of persons at large scale and lead to appreciable

disruptions in the ordinary conduct of social and economic activities." Where the classification applies, human overseers must be able to understand the system's limitations, remain aware of "the possible tendency of automatically relying or over-relying on the output," disregard or reverse that output, and halt the system "in a safe state."

Precision matters here, and overclaiming would be easy. The trigger is acting as a safety component. A read-only analytics tool over building automation data is not automatically high risk, and commercial building heating and cooling is not obviously "the supply of heating," a phrase that points at utility and district scale. The classification question for an individual building's automation system is genuinely unsettled. The defensible statement is that acting moves a system toward that category, not that it lands there.

Industry consensus is blunter. The Open Worldwide Application Security Project names excessive agency as a distinct vulnerability class: "the vulnerability that enables damaging actions to be performed in response to unexpected, ambiguous or manipulated outputs." It also observes that instructing a model to behave "is not a guarantee that the model will always adhere to this."

Peer-reviewed security research closes the loop. Systems that combine retrieval with the ability to invoke tools "blur the line between data and instructions," and processing injected content can "manipulate the application's functionality, and control how and if other APIs are called." A retrieval-plus-action architecture inherits its trust boundary from its data sources. In a building, those sources include whatever a subcontractor typed into a commissioning spreadsheet.

Governance research adds the timing problem. "Without a human in the loop, agents may take multiple consequential actions in rapid succession and bring about significant impacts before a human notices."

Part II: Why This Is a Facilities Problem

Three conditions make the built environment a harder case than the enterprise examples the ontology literature usually cites.

3.1 One word, several meanings, five systems

The clearest published statement of the problem in commercial real estate terms comes from the Open Standards Consortium for Real Estate: "If one system uses 'Pump ID' and another uses 'Asset Tag,' they don't inherently recognize that these labels are referring to the same thing." Bridging that gap requires "complex, manual mappings."

The academic record supplies the general form. A University of Cambridge survey of industry data models found the same physical concept named three different ways across coexisting standards: a "system element" in one, a "functional physical object" in another, a "tagged item" in industry practice. A companion paper poses the ambiguity as a live design question that any facilities professional will recognize: "when referring to a door what elements does this term encompass, for example the door leaf, plus door furniture and/or the door frame?"

Measurement standards carry the same fracture. International and national floor-area standards "often have semantic differences," and the harmonizing international standards "also have differences," which matters because area values drive taxation and valuation. The professional bodies have acknowledged this directly. When the International Facility Management Association and the Building Owners and Managers Association published joint measurement definitions in 2008, the stated reason was that the two standards diverged, with one counting common areas and the other not. Nearly two decades later, computer-aided facility management platforms still ship explicit translation tables mapping one convention onto another. The existence of the translation layer is itself the evidence.

One honest limitation belongs here. Despite searching for it directly, this project found no rigorous published study of a single term such as "asset," "unit," or "property" carrying incompatible meanings across operational commercial real estate systems. The claim is widely made, including in the source corpus that prompted this work, and it is intuitively obvious to anyone who has reconciled a lease system against a maintenance system. It has not been formally studied. That gap is worth naming rather than papering over, and it is a candidate for original Building Lifecycle Management Initiative research.

3.2 Handover, where the record is strongest

The evidence on data loss at lifecycle handover is the most quantitatively solid in this report, and it comes from four independent teams across three countries.

A study of five live United Kingdom building information modeling projects, assessed with a structured quality tool, found error rates of 100 percent on two projects, 84 percent on a third, and 53 percent on a fourth. The fifth scored one percent error but lacked type, component, and system data entirely. The root cause was governance rather than format: "Every subcontractor in the supply chain provided own COBie file with no central point of data check, data federation or quality control." Asset names within a single delivery were "not unique and not consistent," making asset maintenance "almost impossible." In two of the five projects, the client had provided no asset information requirements document at all.

An audit of an operating estate found that 73 percent of assets had no information about their maintenance procedures, and more than half lacked maintenance manual location, warranty, or serial number data. A study of two owner organizations found that neither "had defined information requirements for model based project and information handover," and that they "did not know what to require in models." The same study counted 176 spaces with boundary defects and 61 intersecting spaces in a single government building.

A capability analysis benchmarked the exchange formats themselves against a British asset register standard and found that the industry foundation classes schema lacked direct support for five of 22 required information elements, and the construction operations exchange format lacked six of 22, with the gaps concentrated in cost, written down value, and depreciation.

Perhaps the most telling data point is institutional rather than statistical. Transport for New South Wales published a formal technical standard explaining why it declined to adopt the industry handover format at all, stating that it "has not evolved sufficiently to fulfil the requirements." Among the information it names as unsupported: asset owner, operations and maintenance contract information, asset condition and remaining service life, installation and commissioning details, and asset risk. That same document records that ISO 19650 does not mandate the format, which is worth knowing, because the ISO series is

frequently invoked as though it settles the question. It specifies a management process. It does not specify a data model, a vocabulary, or an exchange format, and this project located no credible empirical study of its adoption or effectiveness.

3.3 Time, which is the cheapest fix and the most expensive omission

The Cambridge survey states the problem in one line worth quoting exactly: process industry data models recognize "that the system elements within a process plant change over time. Building and geographic data models are largely static."

The consequence has now been examined directly. A 2026 study of temporally qualified building elements found that the same physical element "is assigned different Globally Unique Identifiers in each" lifecycle phase, producing "fragmentation of the element's identity." The industry foundation classes standard "does not define a formal mechanism that would link the same physical element across successive phases," and its schema "lacks a conceptual framework linking the states of the same object at different points in time." The results are the loss of property change history and "the inability to formulate cross-phase queries."

A separate capability analysis found that service life planning standards require probability distributions for remaining life, while commercial modeling tools support only simple numeric types, so remaining-life uncertainty cannot be carried in the record at all.

The vocabulary to fix this has existed as a W3C Recommendation since 2013. The provenance ontology supplies entities, activities, and agents together with `generatedAtTime` and `invalidatedAtTime`, which is precisely the validity interval building records lack. This project located no substantive published application of it to commercial real estate or facilities management data. The tool is on the shelf and the industry has not picked it up.

This is the clearest actionable finding in the report. It is cheap to close at design time and expensive to retrofit, and it is the failure most likely to produce a confidently wrong answer that no validator catches. A graph without validity intervals will answer a question about last year using this year's topology, and will not signal that it did so.

Part III: The Audience

4.1 What is known, and what is asserted

The premise that the facilities management and commercial real estate workforce lacks data literacy is widely repeated and thinly evidenced. The three statistics that dominate industry presentations trace to two commercial sources: a business intelligence software vendor, twice, and a paid analyst firm. The most quoted, that only 21 percent of the global workforce is fully confident in its data literacy, measures self-reported confidence rather than tested competence, samples 9,000 office employees across nine countries, and contains no identifiable facilities, property, or building operations population. A peer-

reviewed systematic review of 44 articles concludes that "there is currently no unified definition of Data Literacy, along with its constituent skills," which means the vendor percentages are not measuring a stable construct in the first place.

This report therefore records a negative finding: there is no credible measurement of data literacy in the facilities management or commercial real estate workforce. That gap should be stated rather than filled with a borrowed number.

What is well evidenced is more useful and points in a slightly different direction.

A 2025 survey of more than 1,000 senior commercial real estate decision-makers across 16 markets found that workforce and talent gaps are the number one barrier to an effective technology strategy. Only 33 percent of the workforce feel adequately trained on artificial intelligence. Eighty-one percent of companies report at least three existing systems not producing expected results, and more than 60 percent must resolve fundamental technology problems, including duplicated functionality and dormant systems, before they can use artificial intelligence at all. Ninety percent of firms are testing artificial intelligence solutions. Five percent report achieving their stated objectives.

On building operations specifically, peer-reviewed work across 5,200 buildings names "the lack of in-house staff time or expertise" as an adoption barrier and identifies the need for "a trained building operations workforce." An independent research institute found that "most operators have little to no experience analyzing large amounts of building performance data," and face "a steep learning curve."

The demographic premise rests on primary government statistics rather than survey sentiment. The median United States facilities manager is 53.6 years old, against an all-occupations median of roughly 42. Property and real estate managers sit at 50.1. Professional body research found more members over 70 than under 30 in one facilities management membership, with fewer than 15 percent under 40, and an average member age of 50.9 in another. Nearly 40 percent of facility managers expect to retire within ten years.

4.2 The counter-signal that deserves airtime

One finding cuts against the deficit narrative and should not be suppressed. In a 2026 professional body survey, 72 percent of facility managers said artificial intelligence already affects their daily work, 92 percent expressed interest in adopting additional tools, and nearly 70 percent expected their roles to change significantly within five years. That population is credentialled and more engaged than the workforce at large, which likely overstates the figures. It is still not the profile of a workforce refusing technology.

The problem the evidence supports is capability and data foundations, not appetite. An Awareness message built on the premise that practitioners are resistant will be both wrong and insulting. One built on the premise that they are interested and under-equipped is accurate.

4.3 False familiarity, which is the actual obstacle

The most useful finding in this report for the communication problem comes from cognitive psychology rather than from the built environment.

Two independent literatures converge. The first establishes that people rate their understanding of how things work substantially higher before attempting to explain them than after, and that the effect is specific to explanatory knowledge involving complex causal patterns. It does not appear for factual or procedural knowledge. Experts show a less pronounced version of it. Requesting a mechanistic explanation both exposes the gap and reduces entrenchment.

The second is more directly instructive. The knowledge revision framework establishes that information encoded in long-term memory "leaves a durable trace that cannot simply be erased or overwritten." Incoming information passively reactivates whatever in memory "resonate[s] with that signal as a function of featural overlap." Critically, "both the new information and misconception from prior knowledge must be simultaneously activated" for revision to occur, and "revision can only occur after the newly encoded information becomes integrated with the misconception."

Applied to this topic, the implication is direct and slightly alarming. A facilities professional hearing the word "ontology" will passively reactivate the nearest thing in memory, which is likely a naming convention, a tagging standard, or an equipment classification system. If the Awareness message asserts the correct account without naming the near-miss, the listener absorbs the new material into the old schema, leaves believing they already knew this, and does not return for the next tier.

The obstacle at the Awareness tier is not that the audience will fail to understand. It is that they will believe they already do. That is a design requirement, not a communication risk to be managed afterwards.

Part IV: The Awareness Tier

5.1 Where the scale comes from

No single source originates the Awareness, Knowledge, Application progression. It is a convergent pattern independently instantiated in at least five lineages, and the provenance is worth stating honestly because the three tiers are not equally well grounded.

The closest built-environment analog is the Royal Institution of Chartered Surveyors Assessment of Professional Competence, whose official levels read: Level 1, "knowledge and understanding"; Level 2, "application of knowledge and understanding"; Level 3, "reasoned advice and depth of knowledge." That is quotable verbatim from a current governance document and maps cleanly onto the working scale.

The revised educational taxonomy supplies Remember, Understand, and Apply, though its authors explicitly relaxed the requirement of strict hierarchy. Clinical education supplies Knows, Knows How, Shows How, Does. The information technology skills framework defines its first level of responsibility as performing "routine tasks under close supervision."

Three observations follow, and each affects design.

None of these frameworks has an Awareness tier below its first knowledge tier. The chartered surveyor Level 1 is already knowledge and understanding. The skills framework starts where Application starts. The educational taxonomy's Remember is factual recall, not awareness.

The authority for a distinct Awareness stage comes from change management rather than education, and it is commercially owned. The ADKAR model places Awareness first, and defines it specifically as awareness of the need for change rather than awareness that a change is occurring. Its success criterion is a first-person statement: "I understand the nature of the change and why this change is needed." Awareness explicitly does not require knowing how to execute, agreeing with the reasons, or wanting to change. The model's originator sells the certification, and this project located no independent controlled test of the awareness-before-training sequencing claim. Adult learning theory converges independently through its "need to know why" assumption, but that theory's own author conceded it is "less a theory of adult learning than a model of assumptions."

The tiers therefore rest on unequal footing. The progression from knowledge to application to judgment is well grounded across independent professional frameworks. The specific claim that awareness must come first is grounded in a proprietary practitioner model plus indirect support from learning science. That should be conceded openly rather than dressed up.

5.2 A boundary test

The operational rule that follows from the frameworks is simple enough to apply in a room.

Awareness produces recognition. Knowledge produces judgment. Application produces artifacts.

The test: if the content requires the learner to make a choice, it is Knowledge, not Awareness. If it requires them to build or configure something, it is Application. Anything at the Awareness tier that asks a participant to choose between competing schemas, or to evaluate what a constraint language is for, has crossed the boundary and will produce exactly the disengagement the tier exists to prevent.

A second rule comes from learning science and is less obvious. The expertise reversal effect establishes that "instructional techniques that are highly effective with inexperienced learners can lose their effectiveness and even have negative consequences when used with more experienced learners." Heavy scaffolding helps novices and actively harms people who have moved past that stage. The instructional method must change between tiers, not only the content volume. Delivering Awareness-tier scaffolding to a Knowledge-tier audience is not merely redundant; it degrades their learning.

5.3 What must be communicated at Awareness

Five items, and nothing else.

One. The polysemy recognition, delivered as two cases side by side. The strongest finding in the learning science literature for this purpose is that the mechanism of abstraction is comparison, not exposure. Participants who compared two cases transferred the underlying principle to novel problems at 48 percent, against 19 percent for those who studied the same two cases separately. In a live negotiation exercise, guided comparison with scaffolding reached 90 percent transfer against a 37 percent baseline, and 42 percent of comparison participants articulated the underlying principle completely, against 13 percent. Two concrete cases placed side by side, with the comparison explicitly prompted, does the abstraction work that a stated definition does not. For this audience the two cases write themselves: the

same word meaning different things in the maintenance system and the lease system, and the same physical unit carrying different identifiers at design, commissioning, and operation.

Two. A working definition in one sentence, in their language, delivered after the cases rather than before them. An ontology is the written-down agreement about what the words mean and how the things relate, in a form a machine can check. It is not a database, not a software product, and not a naming standard. The analogy that carries for a facilities audience is the difference between an equipment schedule and a single-line diagram. The schedule says what exists. The diagram says how the pieces connect, which is the only thing that lets someone who has never entered the building answer a question correctly. A variable air volume terminal may sit above one room, belong to no system in that room, be fed by an air handling unit two floors away, and serve a thermal zone spanning several rooms. Four relationships, four questions, four different correct answers, and one generic "belongs to" that destroys all four.

Three. Why now, which is the strongest available argument and requires no technical scaffolding. For decades, when two systems disagreed about what a word meant, a person reconciled it quietly at the boundary. Artificial intelligence does not remove that person by decision. It removes them by acting before anyone reads the output. The fragmentation was always present; the human was absorbing it. This reframes the topic from a new technical burden into a bill coming due on work the profession was never paying for, and it credits the audience rather than implicating them.

Four. What changes when a system acts rather than advises. When software advised, weak definitions produced a bad report and someone caught it. When software acts, weak definitions dispatch a technician to the wrong asset, shut down the wrong air handling unit, or apply the wrong lease clause. The consequence class changes while the defect stays the same. This is the "why we need it," and the regulatory framing supports it without needing to be cited in the room.

Five. What the audience is being asked to do, which at this tier is close to nothing. They are not being asked to build anything, choose a standard, or learn a query language. They are being asked to recognize the failure when they see it, and to understand that the definition of a term is a governed decision with a named owner rather than an accident of whichever system was installed first. The practical deliverable is a person who can ask two questions in a vendor meeting and understand the answers: whose definition of this term does your system use, and what happens when it disagrees with ours.

5.4 What must be handled, not merely avoided

The false familiarity finding converts one item from a caution into a requirement.

An Awareness message that only asserts the correct account will be absorbed into the listener's existing and partially wrong schema. The near-miss belief must be named in the same moment as the correction. In practice that means saying out loud that a naming convention is not an ontology, that a tagging standard is not an ontology, and that an equipment classification system is not an ontology, and saying briefly what each of them does instead. This is uncomfortable, because it tells competent people that something they were proud of does not do the job they assumed. It is also the only mechanism the research supports for actually revising the belief.

A second diagnostic follows from the same literature. Self-rated familiarity is exactly the measure that the illusion inflates. Asking a room whether they are familiar with the term produces a useless answer. Asking someone to explain how the thing works produces an accurate one, and the act of explaining also reduces entrenchment. If the Awareness session needs a check, it should ask for a mechanism, not for a show of hands.

5.5 What does not belong at Awareness

Applying the boundary test rules out, at minimum: any comparison of competing schemas; any treatment of what a constraint language is or where it sits in a stack; the open world and closed world assumptions; upper ontologies; inference and reasoning; anything requiring a decision about tooling; and any implied program of work.

That last exclusion is worth its own sentence, because it is the one most likely to be violated by enthusiasm. The published record prices none of this. No source quantifies what it costs to build or maintain an ontology, no source offers a maintenance staffing model, and no source provides a decay rate for a model left untended. The first question from any owner or executive in the room will be what it costs, and there is no defensible answer available. An Awareness message that implies a build invites a question that cannot be answered, and an unanswered cost question converts interest into skepticism. The tier ends at recognition, deliberately.

Part V: The Honest Caution

One finding should temper expectations for the Awareness tier and should be stated to anyone commissioning the work.

A meta-analysis of 69 studies of end-user training found an overall effect size of 0.75. Decomposed, knowledge and attitudes moved at an effect size of 1.02, which is large. Behavior moved at 0.36, which is small. Field studies reinforce the split. A study of more than 14,000 employees over 15 months found that training delivered at the moment of failure "did not make employees more resilient." A randomized controlled trial with 12,511 participants returned a null result, with a trained group performing marginally worse than control at conventional significance. Later work from the same group reframed the benefit as arising from periodic reminders rather than from training content, "an attention problem rather than a knowledge problem." Research from a government standards body observes that "training employees for policy compliance influences their awareness, unfortunately, it has minimal impact on promoting actual compliant behavior," and that completion metrics are deceptive because people click through.

These studies come from information security, which is the hardest possible test case: adversarial, momentary, and attention-driven. They should not be read as showing that awareness communication is worthless. The same meta-analysis shows it works well for exactly what the Awareness tier is for, which is knowledge and attitude.

What they establish is the correct claim to make. Awareness is a necessary precondition for encoding, not a behavioral lever. It earns its place in the progression because Knowledge and Application cannot be built on a workforce that does not know why the change matters. It does not earn a claim that briefing

people will change what they do. Anyone promising behavior change from an awareness campaign is promising something the evidence does not deliver, and completion rates will not detect the shortfall.

Conclusion

There is an ontology imperative, and it is narrower and better founded than its advocates usually state.

The imperative is not that formal semantics makes artificial intelligence work in buildings. The best available buildings evidence says the opposite: with a well-modeled building graph in place, current systems answer roughly a third of questions usefully, and invent classes that do not exist while doing it. The imperative is that formal semantics is the floor beneath which nothing else is possible. Every serious platform effort builds on one. The two United States federal bodies working on the problem both identify manual mapping as the barrier to scale. And the industry's own economics show the labor of making systems legible sitting between analytics and any return from it.

The distinction matters for how the profession should speak about this. A claim that adopting a standard delivers portable analytics is not supported and will not survive contact with a portfolio. A claim that without shared, machine-checkable definitions the industry cannot reach the capability at all is supported, is stated by NIST in its own words, and is the claim worth making.

Three conditions make this a facilities problem rather than a general data problem. One word carries incompatible meanings across the systems a single building runs on. The handover record is quantifiably poor, with error rates above 80 percent in audited projects and 73 percent of assets in one operating estate carrying no maintenance procedure data at all. And building data models are, in the words of a Cambridge survey, largely static, so records that changed over time cannot say when they changed. That third condition is the cheapest to fix at design and the most expensive to retrofit, and the vocabulary for fixing it has been a W3C Recommendation since 2013 with no substantive built-environment application on record.

For the Awareness tier, the evidence supports a specific and modest design. Lead with two concrete cases and prompt the comparison, because comparison rather than exposure is what produces the abstraction. Define the concept in one sentence afterwards. Name the near-miss beliefs explicitly, because a correction that does not surface the old belief will be absorbed into it. Explain why the change is arriving now in terms that credit the audience for the interpretive work they have been doing invisibly. Explain what changes when software acts rather than advises. Ask for nothing beyond recognition, and stop before implying a program, because the cost question has no published answer.

And expect the tier to move knowledge rather than behavior, because that is what the evidence says it does.

Gaps in the Evidence Base

A program built on this material should know where it is unsupported.

Cost and maintenance economics. No source in this review quantifies what it costs to build or maintain a building ontology, offers a maintenance staffing model, or supplies an empirical decay rate for an untended model. Deployment economics exist for analytics tools overall but do not decompose the metadata share.

Terminology conflict in operational systems. The claim that a single term carries incompatible meanings across commercial real estate systems is widely made and, as far as this review could establish, never formally studied. This is a candidate for original research.

The transaction and portfolio half of the lifecycle. The evidence base covers building operations, handover, and industrial assets. It does not cover lease, transaction, valuation, or portfolio data models. Any whole-of-lifecycle argument drawing on this report should treat that half as unsupported by this material.

ISO 19650 effectiveness. No credible empirical study of adoption rates or effectiveness was located.

Provenance vocabulary in practice. No substantive published application of the W3C provenance ontology to commercial real estate or facilities management data was found.

Data literacy in this workforce. No measurement exists. Every circulating figure is general-workforce, self-reported, and vendor-sponsored.

Awareness-first sequencing. No independent controlled test of the claim that awareness must precede training was located. The claim rests on a proprietary practitioner model.

Adversarial security in building semantics. Beyond general prompt-injection research, there is no treatment of graph poisoning, ontology tampering, or provenance forgery in a built-environment context.

Works Cited

Every link below was fetched and confirmed to resolve to the cited item on 2026-09-04. Open-access full text is preferred where it exists; otherwise the official publisher or issuing-body page is given.

{Paywalled} marks an item whose full text requires payment, subscription, or purchase. Three entries, marked {Browser access only}, are correct publisher landing pages that refuse automated retrieval and render normally in a browser.

- [1] Abdirad, H., & Dossick, C.S. (2019). Normative and descriptive models for COBie implementation: discrepancies and limitations. *Engineering, Construction and Architectural Management*, 26(8), 1820 to 1836.
<https://www.emerald.com/insight/content/doi/10.1108/ecam-10-2018-0443/full/html> {Paywalled} {Browser access only}
- [2] Alarcón, A., de Ramón, J., Ginieis, M., & Andraszak, J. (2026). Data literacy in the labor market: a systematic review. *Humanities and Social Sciences Communications*, 13(1), Article 506.
<https://www.nature.com/articles/s41599-026-06824-w>
- [3] Alnaggar, A., & Pitt, M. (2019). Lifecycle Exchange for Asset Data (LEAD). *Journal of Facilities Management*, 17(5), 385 to 412. <https://discovery.ucl.ac.uk/10079760/1/alnaggar-Pitt.pdf>
- [4] American Council for an Energy-Efficient Economy. King, J., & Perry, C. (2017). Smart Buildings: Using Smart Technology to Save Energy in Existing Buildings. Report A1701.
<https://www.aceee.org/sites/default/files/publications/researchreports/a1701.pdf>
- [5] ASHRAE. (2025). Standards Actions, Vol. XV, Issue 26, June 27 2025.
https://www.ashrae.org/file%20library/technical%20resources/standards%20and%20guidelines/standards%20actions/sqjun_27_2025.pdf
- [6] ASHRAE. (2026). ANSI/ASHRAE Standard 231-2026, A Control Description Language for Building Environmental Control Sequences. https://store.accuristech.com/standards/ashrae-231-2026?product_id=3062190
{Paywalled} Reseller listing; ASHRAE's own bookstore carried no page for this edition at time of verification.
- [7] Balaji, B., et al. (2016). Brick: Towards a Unified Metadata Schema For Buildings. *BuildSys 16*, ACM.
<https://brickschema.org/papers/Brick-BuildSys-2016-Balaji.pdf>
- [8] Balaji, B., et al. (2018). Brick: Metadata schema for portable smart building applications. *Applied Energy*, 226, 1273 to 1292. <https://par.nsf.gov/servlets/purl/10074644>
- [9] Bender, E.M., Gebru, T., McMillan-Major, A., & Shmittchell, S. (2021). On the Dangers of Stochastic Parrots. *FACCT 21*, pp. 610 to 623. ACM. <https://s10251.pcdn.co/pdf/2021-bender-parrots.pdf>
- [10] Bhattacharya, A., et al. (2015). Automated Metadata Construction To Support Portable Building Applications. *BuildSys 15*, ACM. <http://sirrice.github.io/files/papers/arka-buildsys2015.pdf>
- [11] Borkowski, A.S., Jarema, P., Kładź, M., & Smoliar, A. (2026). Temporally Qualified Building Elements. *Technologies*, 14(7), 413. <https://www.mdpi.com/2227-7080/14/7/413>
- [12] Butterfuss, R., & Kendeou, P. (2021). KReC-MD: Knowledge revision with multiple documents. *Educational Psychology Review*, 33, 359 to 398. <https://files.eric.ed.gov/fulltext/ED612635.pdf>
- [13] Cavka, H.B., Staub-French, S., & Poirier, E.A. (2018). Levels of BIM compliance for model handover. *Journal of Information Technology in Construction*, 23, 243 to 258. https://www.itcon.org/papers/2018_12-ITcon-Cavka.pdf
- [14] Chan, A., et al. (2024). Visibility into AI Agents. *FACCT 24*, ACM.
<https://facctconference.org/static/papers24/facct24-63.pdf>

- [15] Chapman, R.E. (2005). Inadequate Interoperability: A Closer Look at the Costs. ISARC 2005. <https://www.iaarc.org/publications/fulltext/isarc2005-85chapman.pdf>
- [16] Deloitte. (2025). 2026 Commercial Real Estate Outlook. <https://www.deloitte.com/us/en/insights/industry/financial-services/commercial-real-estate-outlook.html>
- [17] European Parliament and Council. (2024). Regulation (EU) 2024/1689, Artificial Intelligence Act. Official Journal of the European Union. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- [18] Fierro, G., et al. (2018). Mortar: An Open Testbed for Portable Building Analytics. BuildSys 18, ACM; journal version ACM Transactions on Sensor Networks, 16(1), 2020. <https://gtf.fyi/papers/mortar2018fierro.pdf>
- [19] Fyfe, E.R., McNeil, N.M., Son, J.Y., & Goldstone, R.L. (2014). Concreteness fading in mathematics and science instruction: A systematic review. Educational Psychology Review, 26(1), 9 to 25. <https://link.springer.com/article/10.1007/s10648-014-9249-3> {Paywalled}
- [20] Fyfe, E.R., & Nathan, M.J. (2019). Making concreteness fading more concrete as a theory of instruction for promoting transfer. Educational Review, 71(4), 403 to 435. <https://scholarworks.iu.edu/dspace/items/f1ea5994-0d9f-4215-abec-152a03149277/full>
- [21] Gallaher, M.P., O'Connor, A.C., Dettbarn, J.L. Jr., & Gilday, L.T. (2004). Cost Analysis of Inadequate Interoperability in the U.S. Capital Facilities Industry. NIST GCR 04-867. <https://nvlpubs.nist.gov/nistpubs/gcr/2004/nist.gcr.04-867.pdf>
- [22] Gentner, D., Loewenstein, J., & Thompson, L. (2003). Learning and transfer: A general role for analogical encoding. Journal of Educational Psychology, 95(2), 393 to 408. <https://groups.psych.northwestern.edu/gentner/papers/GentnerLoewensteinThompson03.pdf>
- [23] Greshake, K., et al. (2023). Not What You've Signed Up For: Compromising Real-World LLM-Integrated Applications with Indirect Prompt Injection. AISec 23, ACM CCS. <https://arxiv.org/abs/2302.12173>
- [24] Han, H., et al. (2025). RAG vs. GraphRAG: A Systematic Evaluation and Key Insights. arXiv:2502.11371. <https://arxiv.org/abs/2502.11371>
- [25] Haney, J., & Lutters, W. (2024). From Compliance to Impact. Cyber Security: A Peer-Reviewed Journal, 8(2), 110 to 130. https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=935669
- [26] Hetherington, J., & West, M. (2020). The Pathway Towards an Information Management Framework. Centre for Digital Built Britain, University of Cambridge. https://www.cdbb.cam.ac.uk/files/the_pathway_towards_an_imf.pdf
- [27] Hiatt, J.M. (2006). ADKAR: A Model for Change in Business, Government and our Community. Prosci Learning Center Publications. <https://www.prosci.com/methodology/adkar> Prosci hosts no page for the 2006 book; the official model page, which credits Hiatt as originator, is given instead.
- [28] International Facility Management Association. Core Competencies. <https://www.ifma.org/professional-development/Core-Competency/>
- [29] International Facility Management Association. (2026). 2026 Global Salary Report. <https://blog.ifma.org/what-the-2026-global-salary-report-reveals-about-the-future-of-facility-management> The report itself sits behind an IFMA member login; the official published summary is given.
- [30] International Organization for Standardization. (2018). ISO 19650-1:2018. <https://www.iso.org/standard/68078.html> {Paywalled} International Organization for Standardization. (2020). ISO 19650-3:2020. <https://www.iso.org/standard/75109.html> {Paywalled}
- [31] JLL. (2025). Reality check: The true pace and payoffs of AI adoption in corporate real estate. Global Real Estate Technology Survey 2025. <https://www.jll.com/en-us/insights/global-real-estate-cre-technology-survey>
- [32] JLL. (2025). Global State of Facilities Management Report 2025. <https://www.jll.com/en-us/insights/global-state-of-facilities-management-report>

- [33] Kalyuga, S., Ayres, P., Chandler, P., & Sweller, J. (2003). The expertise reversal effect. *Educational Psychologist*, 38(1), 23 to 31. https://www.tandfonline.com/doi/abs/10.1207/S15326985EP3801_4 {Paywalled} {Browser access only}
- [34] Kara, A., Çağdaş, V., Işıkdag, Ü., & Turan, B.O. (2018). Towards harmonizing property measurement standards. *Journal of Spatial Information Science*, 17, 87 to 119. <https://josis.org/index.php/josis/article/view/97>
- [35] Kendeou, P. (2024). A Theory of Knowledge Revision: The Development of the KReC Framework. *Educational Psychology Review*, 36, 44. <https://link.springer.com/article/10.1007/s10648-024-09885-y> {Paywalled}
- [36] Kramer, H., Lin, G., Curtin, C., Crowe, E., & Granderson, J. (2019). Building analytics and monitoring-based commissioning. *Energy Efficiency*, 12, 2523 to 2544. https://eta-publications.lbl.gov/sites/default/files/building_analytics_-_kramer.pdf
- [37] Krathwohl, D.R. (2002). A Revision of Bloom's Taxonomy: An Overview. *Theory Into Practice*, 41(4), 212 to 218. <https://eric.ed.gov/?id=EJ667155>
- [38] Lahmam Bennani, I., et al. (2021). Query Relaxation for Portable Brick-Based Applications. *BuildSys 21*, ACM. <https://escholarship.org/uc/item/9wm3591h>
- [39] Lain, D., Kostianen, K., & Čapkun, S. (2022). Phishing in Organizations. *IEEE Symposium on Security and Privacy*. <https://arxiv.org/abs/2112.07498>
- [40] Langer, G., Hirsch, T., Kern, R., Kohl, T., & Schweiger, G. (2025). Large Language Models for Fault Detection in Buildings' HVAC Systems. *Energy Informatics*, LNCS 15272. https://link.springer.com/chapter/10.1007/978-3-031-74741-0_4 {Paywalled}
- [41] Leal, D., Cook, A., Partridge, C., Sullivan, J., & West, M. (2020). A Survey of Industry Data Models and Reference Data Libraries. Centre for Digital Built Britain, University of Cambridge. <https://www.repository.cam.ac.uk/items/cf5f76da-2dce-4a57-866c-8262784b984d>
- [42] Lebo, T., Sahoo, S., & McGuinness, D., eds. (2013). PROV-O: The PROV Ontology. W3C Recommendation. <https://www.w3.org/TR/prov-o/>
- [43] Lin, G., Kramer, H., & Granderson, J. (2020). Building fault detection and diagnostics. *Building and Environment*, 168, 106505. <https://escholarship.org/uc/item/5cb0g38t>
- [44] Lin, G., Kramer, H., Nibler, V., Crowe, E., & Granderson, J. (2022). Building Analytics Tool Deployment at Scale. *Energies*, 15(13), 4858. <https://www.mdpi.com/1996-1073/15/13/4858>
- [45] Mavrokapnidis, D., Fierro, G., Korolija, I., & Rovas, D. (2023). A Programming Model for Portable Fault Detection and Diagnosis. *e-Energy* 23, ACM. <https://gtf.fyi/papers/mavrokapnidis2023programming.pdf>
- [46] McGrath, V. (2009). Reviewing the Evidence on How Adult Students Learn. *The Adult Learner*, 99 to 110. <https://files.eric.ed.gov/fulltext/EJ860562.pdf>
- [47] Miller, G.E. (1990). The assessment of clinical skills/competence/performance. *Academic Medicine*, 65(9 Suppl), S63 to S67. <https://academic.oup.com/academicmedicine/article/65/9/S63/8361091> {Paywalled} Abstract-level record also verified at <https://pubmed.ncbi.nlm.nih.gov/2400509/>
- [48] Mulayim, O.B., et al. (2024). Large Language Models for the Creation and Use of Semantic Ontologies in Buildings. *BuildSys 24*, ACM. <https://gtf.fyi/papers/large2024mulayim.pdf>
- [49] Mulayim, O.B., et al. (2025). BuildingQA: A Benchmark for Natural Language Question Answering over Building Knowledge Graphs. *BuildSys 25*, ACM. <https://gtf.fyi/papers/buildingqa2025mulayim.pdf>
- [50] National Institute of Building Sciences. (2023). COBie Version 3, in NBIMS-US V4. <https://nibs.org/nbims/v3/cobie>
- [51] National Institute of Standards and Technology. Building Digitization and Semantic Interoperability program. <https://www.nist.gov/programs-projects/building-digitization-and-semantic-interoperability>

- [52] National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework 1.0. NIST AI 100-1. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>
- [53] National Institute of Standards and Technology. (2024). Generative Artificial Intelligence Profile. NIST AI 600-1. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- [54] Open Standards Consortium for Real Estate. (2026). OSCRE Unveils Vision for a Smart Data Highway. <https://www.prweb.com/releases/oscre-unveils-vision-for-a-smart-data-highway-to-take-real-estate-data-standards-to-the-next-level-302696018.html> OSCRE hosts its own copy of this release, but that file blocks automated retrieval; the wire-service version carries the identical headline, date, and issuer.
- [55] OWASP Foundation. (2024). OWASP Top 10 for Large Language Model Applications, Version 2025. <https://owasp.org/www-project-top-10-for-large-language-model-applications/assets/PDF/OWASP-Top-10-for-LLMs-v2025.pdf>
- [56] Pan, S., et al. (2024). Unifying Large Language Models and Knowledge Graphs: A Roadmap. IEEE Transactions on Knowledge and Data Engineering. <https://arxiv.org/abs/2306.08302> Preprint of record; the IEEE version of record is paywalled.
- [57] Patacas, J., Dawood, N., Vukovic, V., & Kassem, M. (2015). BIM for facilities management: evaluating BIM standards in asset register creation and service life planning. Journal of Information Technology in Construction, 20, 313 to 331. https://www.itcon.org/papers/2015_20.content.00383.pdf
- [58] Paul, L., et al. (2025). Open building operating system. Science and Technology for the Built Environment, 31(3), 294 to 311. <https://escholarship.org/uc/item/0zv1134v>
- [59] Prümmer, J., van Steen, T., & van den Berg, B. (2024). Assessing the effect of cybersecurity training on end-users: A meta-analysis. Computers and Security, 150, 104206. <https://scholarlypublications.universiteitleiden.nl/handle/1887/4195663> {Paywalled} Repository record; the publisher full text is closed access.
- [60] RICS. (2023). UK Facilities Management Survey, Q2 2023. <https://www.rics.org/content/dam/ricsglobal/documents/market-surveys/uk-facilities-management-survey/Q2%202023%20RICS%20Facilities%20Management%20Survey.pdf>
- [61] RICS. (2025). Assessment of Professional Competence: Assessor Guide, June 2025. https://www.rics.org/content/dam/ricsglobal/documents/surveying/APC-assessor-guide_june-25.pdf
- [62] RICS & International Facility Management Association. (2017). Raising the Bar III. https://www.rics.org/content/dam/ricsglobal/documents/to-be-sorted/raising_the_bar_iii_from_operational_excellence_to_strategic_impact_in_fm_rics.pdf
- [63] Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. Cognitive Science, 26(5), 521 to 562. <https://cogdevlab.yale.edu/sites/default/files/files/rozenblit%20&%20keil%20%202002.pdf>
- [64] Saha, A., & Paul, L. (2024). Reduce Barriers to Deployment of Building Control and Analytics Applications. US DOE Building Technologies Office Peer Review. <https://www.energy.gov/sites/default/files/2024-11/bto-peer-2024-34669-32619-Reduce-Barriers-Deployment-Building-Saha-Paul.pdf>
- [65] Sequeda, J., Allemang, D., & Jacob, B. (2024). A Benchmark to Understand the Role of Knowledge Graphs on Large Language Model's Accuracy for Question Answering on Enterprise SQL Databases. GRADES-NDA, ACM. <https://arxiv.org/abs/2311.07509> Preprint of record; the ACM version of record is paywalled.
- [66] SFIA Foundation. SFIA 9, Levels of Responsibility. <https://sfia-online.org/en/sfia-9/responsibilities>
- [67] Sweller, J. (1988). Cognitive load during problem solving. Cognitive Science, 12(2), 257 to 285. https://onlinelibrary.wiley.com/doi/10.1207/s15516709cog1202_4 {Paywalled} {Browser access only}

- [68] Transport for New South Wales. (2022). Why Not COBie? DMS-SD-138 v3.1.
<https://www.transport.nsw.gov.au/system/files/media/documents/2023/Why-not-COBie-DMS-SD-138.pdf>
- [69] Tsay, G.S., Staub-French, S., & Poirier, É. (2022). BIM for Facilities Management. Applied Sciences, 12(19), 9542.
<https://www.mdpi.com/2076-3417/12/19/9542>
- [70] Ulbrich, C. (2026). Real estate's AI success depends on people not technology. World Economic Forum.
<https://www.weforum.org/stories/2026/01/real-estate-ai-success-depends-on-people/>
- [71] United States Bureau of Labor Statistics. (2024). Current Population Survey, Household Data Annual Averages, Table 11b, 2023 annual averages. <https://www.bls.gov/cps/data/aa2023/cpsaat11b.pdf>
- [72] United States Department of Energy. Semantic Modeling and Interoperability.
<https://www.energy.gov/cmei/buildings/semantic-modeling-and-interoperability> United States Department of Energy. ASHRAE Standard 223P. <https://www.energy.gov/cmei/buildings/ashrae-standard-223p>
- [73] Wei, A., Haghtalab, N., & Steinhardt, J. (2023). Jailbroken: How Does LLM Safety Training Fail? NeurIPS 36.
https://papers.nips.cc/paper_files/paper/2023/hash/fd6613131889a4b656206c50a8bd7790-Abstract-Conference.html

About BLMI

The Building Lifecycle Management Initiative (BLMI) is a volunteer-led, not-for-profit initiative working at the intersection of Facility Management, Commercial Real Estate, PropTech, and Data Governance. BLMI advances a single foundational premise: that buildings generate valuable data across their entire lifespan (from design through deconstruction) and that the industry's persistent failure to preserve and transfer that data across ownership transitions, operational phases, and technology generations remains an industry-wide challenge the built environment has been slow to systematically address.

BLMI does not create new standards. Its role is synthesis and advocacy: connecting the work of established frameworks (IFMA, BOMA, CoreNet Global, OSCRE, ASHRAE, and USGBC) into a coherent lifecycle perspective and translating that perspective into actionable guidance for the professionals who manage, own, and operate buildings at scale.

Publications are produced through a research and editorial process developed by domain practitioners and made freely available to the industry. To contribute, collaborate, or access the full research library, visit blmi.org.

blmi.org

Disclaimer

The content of this publication is provided for informational purposes only and does not constitute legal, financial, or professional advice. While reasonable care has been taken to ensure the accuracy and currency of the information presented, Building Lifecycle Management Initiative makes no representations or warranties regarding its completeness or fitness for a particular purpose. References to third-party organizations, standards bodies, products, or services do not imply endorsement. Sponsor relationships do not influence editorial content or conclusions. Views expressed are those of the authors and do not necessarily represent the positions of any affiliated organization, standards body, or sponsor.

© 2026 Building Lifecycle Management Initiative. This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). creativecommons.org/licenses/by/4.0

AI Use Transparency

This publication was produced with the assistance of AI tools in content drafting, research synthesis, and editorial review. All strategic analysis, source selection, factual claims, and interpretive conclusions represent the professional judgment of Building Lifecycle Management Initiative's human contributors. Every factual assertion has been verified against primary sources prior to publication. Building Lifecycle Management Initiative maintains full editorial independence; no claim appears in this publication solely on the basis of AI-generated content.

Produced under Building Lifecycle Management Initiative AI Use Policy | Effective April 2026