FACILITIES AND ASSET MANAGEMENT

Smart Building Asset Management Frameworks for Reducing Equipment Downtime, Maintenance Costs and Operational Risks in Saudi Arabia

Muhammad Junaid Raza

Published November 2026

Abstract

Smart building investment in Saudi Arabia is expanding the installed base of digitally connected mechanical, electrical and plumbing assets, yet connected equipment does not automatically deliver reliable operations. This critical integrative review examines how asset management frameworks can reduce unplanned downtime, control whole-life maintenance expenditure and contain operational risk across complex building portfolios. Thirty DOI-identified articles published during 2020–2025 provide a bounded interpretive evidence base, complemented by two supplied technical examples. Evidence is compared across information continuity, sensing and fault diagnosis, maintenance prioritisation, decision accountability and performance verification. The resulting framework joins an authoritative asset register to condition intelligence, failure consequence analysis, executable work orders and post-intervention learning. Its distinctive proposition is that economic and reliability gains arise from disciplined decisions, not predictive models alone. Saudi implementation requires special attention to cooling dependency, climatic exposure, heterogeneous building maturity, contractor interfaces and data governance. The review offers measurable indicators, a phased adoption pathway and falsifiable research propositions. It does not claim that the proposed architecture has been field tested or that disparate case study savings can be generalised to Saudi portfolios.

Keywords: Smart buildings; asset management; predictive maintenance; facilities management; operational resilience; Saudi Arabia; digital twins; lifecycle cost

1 Introduction

The operational value of a smart building is realised after handover, when its heating, ventilation and air conditioning (HVAC), electrical distribution, vertical transport, water systems and life-safety equipment must remain dependable despite aging, load variation and constrained access. A sophisticated automation interface may expose alarms without clarifying which component is deteriorating, who owns the corrective decision or whether the intervention restored performance. This distinction matters in Saudi Arabia, where intense cooling demand can magnify the service consequences of plant failure. Reliability therefore warrants treatment as a portfolio governance problem rather than a narrow exercise in algorithm selection (Chen et al., 2023; Hakimi et al., 2024).

Conventional maintenance programmes normally combine statutory inspections, calendar-based preventive work and reactive repairs. Their weaknesses are not uniform. Preventive servicing may replace functioning parts too soon; corrective work can trigger costly interruption; and condition monitoring can create more alerts than technicians can investigate. Building information modelling, sensor networks and digital twins promise richer equipment context, but reports repeatedly identify incomplete asset data, weak interoperability and limited field validation (Hosamo et al., 2022a,

2022b; Gispert et al., 2025; Asare et al., 2024). Adding technology to fragmented processes can simply accelerate the production of ambiguous warnings.

This review treats downtime, maintenance cost and operational risk as interdependent outcomes. A short repair is not necessarily an effective repair if recurrence rises; a low-cost maintenance regime is not economical if it exposes occupants to unacceptable failures; and conservative blanket replacement can reduce immediate hazards while wasting capital. The analytical task is to establish how asset information, fault intelligence and managerial controls can be coordinated to make these trade offs explicit. Saudi Arabia supplies an important application context because Vision 2030 building development combines new digitally specified assets with legacy facilities, diverse climatic zones and intensive service expectations (Hegazy, 2025; Olawale et al., 2025).

The contribution is a decision-centred Smart Building Asset Assurance Framework. Unlike reviews confined to fault detection, it specifies the transitions from detected degradation to authorised work, verified restoration and revised maintenance policy. It also clarifies where digital capabilities are necessary but insufficient. The framework is intended as an evaluable design proposition for owners, operators and researchers, not as an empirical finding from a Saudi building sample.

1.1 Aim, objectives and research questions

The aim is to critically synthesise 2020–2025 scholarship and formulate an operationally implementable framework that links smart building asset intelligence to lower equipment downtime, defensible maintenance expenditure and reduced operational exposure in Saudi Arabia. Four objectives guide the inquiry: establish the minimum trustworthy asset information architecture; evaluate evidence for condition monitoring and diagnostic methods; connect reliability and consequence analysis to maintenance prioritisation; and specify governance and measurement arrangements that make performance improvement auditable.

Three questions structure the argument. First, which information and analytical capabilities reliably support earlier recognition of asset deterioration? Second, how should maintenance teams translate uncertain fault indications into proportionate actions when safety, comfort, cost and continuity compete? Third, which contextual and institutional conditions determine whether a framework developed across international buildings can be transferred credibly to Saudi facilities? These questions separate technical possibility, organisational decision quality and local applicability.

2 Review methodology

A critical integrative literature review was selected because the research problem crosses building engineering, information modelling, reliability management and organisational control. A statistical synthesis would be inappropriate: reported studies differ substantially in assets, failure definitions, sampling periods, baseline regimes and outcome metrics. The review instead compares mechanisms, boundary conditions and contradictory interpretations. Its final product is a reasoned framework and a set of testable claims, rather than a pooled numerical effect estimate (Hakimi et al., 2024; Ma et al., 2024).

The review corpus comprises thirty identifiable journal articles dated 2020–2025. Searches were oriented around combinations of “building asset management”, “facilities management”, “digital twin”, “condition-based maintenance”, “HVAC fault detection”, “building information modelling”, “maintenance risk” and “Saudi Arabia”. Relevant publisher and university records were consulted to check

publication details and DOI strings. The two user supplied articles informed the comparison of Saudi digital transformation and event-driven building automation, although the former, published in 2026, was deliberately excluded from the date restricted reference list. No database wide retrieval counts or independently replicated screening totals are claimed.

Eligible sources contributed a defined mechanism to at least one of the following: asset data continuity, diagnostic performance, maintenance decision making, lifecycle economics, resilience or Saudi operational context. Materials were excluded when they lacked a traceable journal DOI, fell outside the time window, or addressed unrelated industrial systems without a defensible connection to buildings. Review papers and individual implementations were retained but weighted differently: surveys identify recurring concepts, whereas prototypes illustrate feasibility within narrower settings. Citation alone did not establish a causal maintenance benefit.

The synthesis used a structured comparison matrix recording asset class, data inputs, decision output, validation type, claimed outcome and implementation constraint. Themes were developed by contrasting evidence for sensing, models, processes and governance, with particular attention to missing links between algorithm outputs and executed maintenance. The resulting findings are interpretive, not a comprehensive systematic census. Publication bias, inaccessible proprietary performance data, English language concentration and the absence of primary Saudi site observations limit external validity. These limitations are disclosed rather than obscured by artificial screening statistics.

3 Evidence synthesis: From asset visibility to reliability

The literature establishes a recurring architectural problem: building design information and operating information are often held in different systems, with different identifiers and different accountability. Digital models can describe geometry and equipment relationships, while building management systems record temperatures, pressures and alarms, and computerised maintenance management systems record failures, labour and parts. Unless each record maps to the same maintained asset, a predictive alert cannot be reliably translated into a physical inspection or followed through to lifecycle cost reporting (Valinejadshoubi et al., 2022; Gispert et al., 2025).

An authoritative asset register is consequently a prerequisite for advanced prediction. Each critical item needs a persistent identifier, precise physical location, parent system, model and serial information, capacity, commissioning baseline, maintainability constraints and relevant safety classification. The data model must also preserve relationships: an air handling unit depends on electrical supply, chilled water circulation, actuators and controls. When a dependent component fails, the observed symptom can appear elsewhere. System-level topology improves root cause interpretation and prevents technicians from treating each alarm as an independent event (Long et al., 2024; Zhang et al., 2023).

The appropriate digital model is not necessarily a photorealistic twin. Evidence reviews distinguish geometry-heavy representations from operationally useful models that synchronize selected state variables, maintenance history and equipment relationships. A lightweight information model can be preferable where the principal decisions concern failure mode, work access and replacement lead time. Full simulation becomes justified when dynamic interaction, control tuning or counterfactual assessment adds demonstrable value. This distinction restrains expenditure on visual complexity that does not improve reliability decisions (AlBalkhy et al., 2024; Hakimi et al., 2024).

Data quality creates a further boundary. Missing sensor calibration, inconsistent engineering units, clock drift and unrecorded setpoint changes can make normal operation appear faulty. The review by Zhang et al. (2021) highlights the need to evaluate sensor usefulness rather than simply increase sensing density. Operational commissioning should therefore include point to asset mapping, expected operating ranges, timestamp validation and confidence flags. Where signal quality is inadequate, the defensible response is targeted investigation rather than automated component replacement.

Equipment classification is another neglected condition of comparability. Portfolio-level failure reports often mix critical and noncritical assets, planned service interruptions and accidental failures, or component defects and system outages. A heat exchanger operating with reduced efficiency may create energy penalties without immediate loss of service, while a failed controller can disable otherwise functional plant. Analytical datasets should preserve these distinctions and specify failure boundaries. Otherwise an apparent reduction in event counts may represent changed coding rather than improved engineering reliability. Consistent taxonomies also improve communication between owners, specialist contractors and insurers.

3.1 Diagnostics, prognosis and actionable uncertainty

Fault detection identifies anomalous operation; diagnosis attributes a plausible cause; prognosis estimates how the condition may develop; and maintenance optimisation determines whether, when and how to intervene. These are analytically distinct tasks. HVAC literature documents rule-based methods, statistical models, supervised machine learning and hybrid approaches, but no single family reliably dominates every equipment type or building configuration (Mirnaghi & Haghighat, 2020; Chen et al., 2023; Matetić et al., 2023). Selecting a model without first specifying the decision it supports reverses the proper engineering sequence.

For cooling plant, interpretable rule checks may detect simultaneous heating and cooling, persistent valve hunting or implausible approach temperatures. Statistical residuals can reveal equipment drift after accounting for weather and occupancy. Supervised classifiers may recognise known fault signatures, but training labels are frequently sparse, inconsistent or affected by maintenance actions that prevent actual failure. A prognostic estimate should therefore report confidence and the consequences of error, rather than present a single remaining life number as established fact (Melgaard et al., 2022; Chen et al., 2022).

A particularly important weakness is false alarm economics. Raising sensitivity may reduce missed faults while increasing inspections, technician interruption and unnecessary shutdowns. Conversely, excessively conservative thresholds can postpone intervention until service is lost. Model performance must be assessed with fault prevalence, alert lead time, false alarms per operating month, diagnosis correctness and subsequent work order outcomes. Laboratory accuracy alone rarely captures these operational burdens. Studies of AI-based HVAC diagnosis reinforce the necessity of realistic evaluation data and fair comparisons across baseline methods (Bi et al., 2024; Chen et al., 2023).

Edge analytics offer a practical option where continuous cloud connectivity is undesirable or where local response is time-sensitive. Alrashidi et al. (2025) examined low-power event-driven processing for adaptive building control and anomaly recognition, reporting favourable results in the evaluated setting. Their findings are relevant to computational architecture, not proof that the same percentage benefits would arise from Saudi portfolio-wide maintenance. Sensor configuration, intervention

policies and building type mediate transfer. Human authorisation remains appropriate for actions with life-safety implications or potential disruption to critical services.

Digital twins extend diagnostic approaches by maintaining a contextual representation of asset state. Their most credible maintenance contribution is the ability to associate a deviation with asset history, dependency structure and alternative operating scenarios. Yet digital twin literature also identifies limited evidence that prototypes consistently improve field maintenance outcomes at scale (Hosamo et al., 2022a, 2022b; Asare et al., 2024). The central question is therefore not whether a virtual model exists, but whether it shortens time to correct diagnosis and increases the fraction of effective interventions.

Maintenance datasets contain selection bias because inspections and repairs alter the observed natural history of faults. Assets that appear exceptionally reliable may simply have undergone preventive replacement before deterioration became measurable. Training a model on the absence of recorded failure can therefore mistake intervention for equipment health. Researchers should explicitly document service history, censoring, sensor availability and periods of abnormal operation. Comparative evaluation against simple engineering thresholds is essential: an elaborate algorithm should demonstrate incremental decision value, not merely outperform an artificially weak baseline. These requirements are consistent with cautionary findings across fault diagnosis reviews.

3.2 Downtime as an end-to-end operational outcome

Downtime should be measured from the onset of unacceptable service degradation to confirmed functional restoration, with planned outages distinguished from unplanned incidents. This definition exposes more than repair duration. It includes fault discovery, triage, work authorisation, access arrangements, spare part availability, technical execution and recommissioning. An advanced diagnostic system can cut detection time without materially reducing total downtime if procurement and contractor response dominate the incident timeline. Consequently, optimisation must address both technical detection and the work management chain.

Reliability prioritisation should reflect service criticality, failure frequency and restoration difficulty. A fault affecting a single noncritical zone does not warrant the same escalation as one compromising a main chilled water distribution loop or life-safety interface. Failure mode analysis provides structure, but simple risk scores can conceal severe low-frequency consequences. Critical risks require explicit escalation rules, designated decision owners and contingency provisions, even where their estimated probability is low (Izaddoost et al., 2021).

Work order automation is useful only when records contain actionable information. A coherent alert should identify equipment, probable failure mode, evidence, urgency, operational consequence, safe inspection steps and the required skills. Research using maintenance request histories suggests that machine learning may help classify service requests and allocate technical resources (D'Orazio et al., 2023). Nevertheless, automated priority labels require review whenever occupant safety, regulatory obligations or specialist judgement are involved. An algorithm may propose a queue; accountable operations management must approve consequential departures from established controls.

Service restoration also has an organisational dimension beyond physical repair. If plant operation is resumed without verifying protective devices, controller sequences or air quality conditions, nominal uptime may conceal heightened risk. Recommissioning criteria should therefore be proportionate to failure consequence. For a minor terminal unit, a functional spot check may suffice; for a central plant or critical electrical feeder, documented tests and supervisory sign off are more appropriate.

Measuring restoration against verified performance prevents maintenance teams from being rewarded merely for closing tickets rapidly. It also creates traceable evidence for subsequent reliability investigation.

3.3 Maintenance economics and whole-life cost

Maintenance cost is a composite measure covering routine labour, corrective labour, consumables, replacement components, specialist contractors, temporary services, procurement delay, lost productive use and premature capital renewal. Saving on one category while increasing another cannot establish improved economic performance. Owners therefore need whole-life comparisons against a credible baseline, with equipment class and service intensity held comparable. Cost per maintained asset and cost per operating hour are more interpretable when paired with reliability and service quality measures.

Predictive maintenance can create value through avoided catastrophic failure, more efficient grouping of planned visits, better spare parts forecasting and extended useful life. It can also create new expenditures: additional instrumentation, licences, integration work, cybersecurity management, data engineering and model upkeep. The investment case must compare these continuing costs against reasonably attributed savings. Extrapolating an experimental energy reduction into a maintenance cost saving is invalid unless an empirical relationship is demonstrated. Even when adaptive control reduces electrical demand, its effect on component wear may vary by cycling pattern, load and equipment design (Alrashidi et al., 2025; Alghassab, 2024).

A financially defensible evaluation considers discounted costs over an explicit horizon, risk-weighted service losses and uncertainty around failures avoided. Decision thresholds should distinguish inexpensive inspection from invasive repair and major renewal. Where failure consequence is high but prediction confidence low, the rational action may be additional sensing or functional testing rather than immediate replacement. Where component availability is limited, strategically positioned spares may deliver more continuity than a marginal improvement in prediction accuracy. Procurement and maintenance planning are therefore inseparable parts of asset governance.

4 Saudi Arabian operating conditions and transferability

Saudi Arabia is not one homogeneous building environment. Coastal humidity and corrosion, inland dust exposure, extreme summer temperatures, differences in building typology and variations in occupancy all affect failure patterns and sensor interpretation. Asset models calibrated on temperate climates can misclassify normal cooling load or overlook local degradation mechanisms. Saudi building code and energy performance studies underline the importance of geographic and climatic differentiation rather than assuming one operational baseline nationwide (Olawale et al., 2025).

Cooling continuity warrants particular attention because disruption can affect thermal comfort, occupancy, business operation and the operating envelope of sensitive equipment. Plant redundancy, chilled water balancing, condenser performance, filtration and electrical supply quality are consequently relevant to asset criticality. The framework should classify failure consequences by building purpose: a hospital, shopping centre, school, hospitality facility and data centre may assign markedly different weights to identical equipment faults. This is a proposed contextual adaptation, not a reported Saudi wide risk ranking.

The national smart city agenda provides an enabling policy context, but programmatic ambition should not be confused with completed operational capability. Hegazy (2025) identifies digital city opportunities alongside data security and infrastructure readiness constraints. Saudi government building energy management evidence also indicates the practical value of formal measurement, management commitment and performance review (Alotaibi et al., 2025). For asset programmes, this means procurement specifications, handover requirements and service contracts should mandate usable maintenance information rather than merely marketing connectivity.

Existing buildings present different implementation economics from new developments. Retrofitting full digital models may be expensive where drawings are inaccurate and plant metadata is missing. A risk-based inventory and selected high-value sensors can establish a stronger foundation than rapid deployment of an exhaustive three dimensional twin. In new buildings, maintainable data should be required from design through commissioning, including verified identifiers and operating baselines. Barriers identified in BIM adoption studies, notably expertise and investment, support a differentiated pathway rather than a mandatory single technology stack (Durdyev et al., 2022; Leygonie et al., 2022).

5 Proposed Smart Building Asset Assurance Framework

The proposed framework contains six connected decision layers. Layer one establishes asset identity and criticality; layer two assures trustworthy operating data; layer three interprets faults and degradation; layer four evaluates service consequence, uncertainty and costs; layer five executes controlled maintenance actions; and layer six tests whether the intervention achieved intended outcomes. The layers are sequential for accountability but iterative in operation. A completed repair updates asset history, model assumptions and future intervention policy. Figure 1 illustrates this feedback architecture as a conceptual proposal synthesized from the reviewed literature (Hakimi et al., 2024; Gispert et al., 2025; Ma et al., 2024).

Asset identity anchors every subsequent step. Required records include technical specifications, dependencies, baseline performance, prior faults, contractor obligations and renewal assumptions. Data assurance then checks signal provenance, calibration, completeness and operating context. Diagnosis translates trustworthy measurements into competing explanations, stating uncertainty rather than suppressing it. Decision support compares inspection, monitoring, adjustment, repair and replacement against risk tolerances and predicted service effects. A work management interface assigns ownership and enforces completion evidence.

The critical architectural innovation is a verification gate between work completion and declared benefit. Closing a ticket confirms administrative activity, not necessarily asset recovery. Verification should combine post-maintenance performance, recurrence surveillance, comfort or service indicators and an explicit acceptance threshold. Failed verification returns the asset to diagnosis with the original and revised hypotheses preserved. This procedure links digital intelligence to operational learning and reduces the chance that repeated symptoms are recorded as unrelated maintenance successes.

Governance responsibilities remain visible across layers. Owners define risk appetite and investment boundaries; engineering managers approve technical policies; maintainers document interventions; data stewards manage identifiers and quality; and independent assurance verifies safety critical controls. The architecture can accommodate commercial building management platforms, but it should remain vendor neutral by specifying interfaces and data semantics rather than a preferred

product. Semantic asset models become particularly valuable when portfolios span multiple contractors and technology generations (Gispert et al., 2025; Dahanayake & Sumanarathna, 2021).

The six layers and their verification feedback are shown in Figure 1.

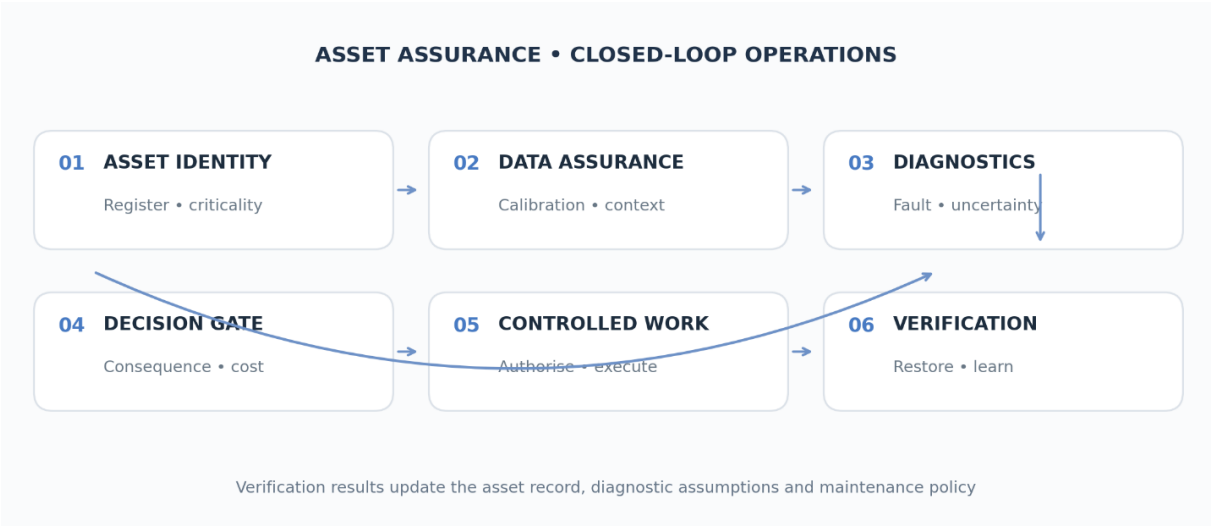


Figure 1. Proposed six-layer asset assurance architecture with verification feedback (original conceptual diagram).

5.1 Prioritisation and measurement model

A practical prioritisation score can combine consequence severity, evidence strength, deterioration rate and expected time to service loss. The score is a triage aid, not an automatic authorisation rule. Critical safety conditions must override arithmetic rankings, and an uncertain but potentially catastrophic fault may require prompt inspection despite modest predicted likelihood. Maintenance managers should log the rationale for overrides to permit later learning. Table 1 maps evidence supported capabilities to practical failure points and corresponding controls.

Evaluation requires a small but balanced indicator set. Unplanned downtime hours should be normalised by operating exposure and supplemented by mean time to restoration. Repeated failures within a defined interval indicate ineffective interventions. Maintenance cost should include analytics operating expenditure, not merely technician invoices. Risk indicators should measure overdue statutory actions, critical alarms awaiting review and documented near misses. Figure 2 depicts the directional hypothesis: better data and more reliable decisions can influence three outcomes through different pathways; it deliberately avoids numerical effect sizes that have not been established.

A useful comparative design is a staggered implementation across comparable buildings or asset classes. Establish a pre-intervention baseline, identify seasonal demand and occupancy effects, and compare outcomes after deployment with a suitable concurrent control. Track failure severity, equipment age and changes in contracted service levels. An interrupted time series or difference in differences analysis may help distinguish genuine maintenance benefit from unrelated changes in capital investment or demand. Table 2 specifies indicators and the principal confounders that a future Saudi evaluation should control.

The principal capabilities, their implementation risks and the corresponding management controls are summarised in Table 1.

Table 1. Capability-to-control matrix for smart building asset management.

Capability	Principal implementation risk	Required management control
Equipment identity	Unmatched BMS and work order identifiers	Asset register with persistent IDs and topology
Condition intelligence	Sensor drift or misleading alarms	Commissioning, calibration and confidence flags
Diagnostic models	False alerts and overconfident predictions	Baseline comparison and uncertainty thresholds
Work execution	Correct insight without timely repair	Named decision owner, spares and escalation route
Performance assurance	Ticket closure without functional recovery	Acceptance tests and recurrence monitoring

Figure 2 sets out the hypothesised pathways from information quality to operational outcomes.

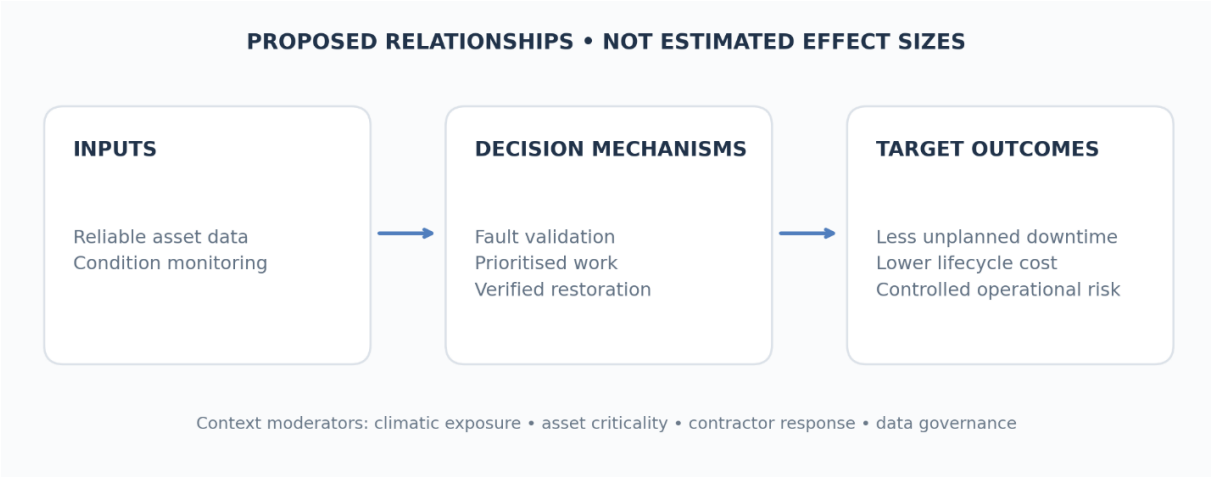


Figure 2. Hypothesised pathways from information quality to operational outcomes (qualitative; not an effect-size chart).

The proposed evaluation indicators are listed in Table 2.

Table 2. Proposed indicators for controlled, comparative field evaluation.

Outcome	Operational indicator	Measurement safeguard
Continuity	Unplanned outage hours / 1,000 operating hours	Separate planned work; record severity and duration
Restoration	Median elapsed time from incident to verified recovery	Track diagnosis, procurement and execution intervals
Maintenance value	Total maintenance and analytics cost / asset year	Include setup, licence, labour and replacement expenditure
Intervention quality	Repeat fault rate within 30/90 days	Use consistent fault codes and matched equipment
Risk control	Overdue critical corrective actions and near misses	Audit safety exceptions and changing reporting coverage

6 Discussion and theoretical implications

Three deductions emerge from the synthesis. First, digital capability is a conditional resource rather than a direct cause of reliability. It creates useful information only when identifiers, measurements and contextual relationships are sufficiently accurate. Second, decisions mediate outcomes: warning lead time has limited value unless staff can procure, authorise and execute effective action. Third, verification closes the causal chain. Without evidence that a repair restored function, organisations cannot distinguish predictive success from administrative closure. These deductions extend technology-centred discussions towards accountable operational management.

The literature also reveals an important asymmetry across outcomes. Downtime reduction may be achieved through quicker response and spare availability without any machine learning model. Maintenance cost reduction may require a longer observation horizon because implementation expenditure arrives early while avoided failures are uncertain. Operational risk can improve through redundancy, training and procedural checks even when equipment condition is unchanged. A credible framework should therefore resist collapsing all benefits into a single smartness index. Interventions must be selected according to the specific mechanism by which they change each outcome.

For theory, the framework proposes that information quality affects reliability through the joint mediation of diagnosis accuracy and execution capability. Contractor coordination and asset criticality moderate this relationship; verification maturity determines whether benefits persist. These propositions are testable by comparing sites with similar sensor investment but different work order governance. They also explain why vendor demonstrations cannot establish portfolio-level generalisability. The appropriate unit of analysis may be the asset–work order episode rather than the building dashboard, because that episode connects observation, action and measurable consequence.

Prior research on digital twins emphasises interoperability and continuous information exchange (Zhang et al., 2023; AlBalkhy et al., 2024). The present model adds a managerial requirement: each exchange must support an auditable decision. A well-connected digital twin that is not linked to maintenance authorisation or performance acceptance remains operationally incomplete. Conversely, a modest asset register and calibrated sensors may deliver substantial value where governance is strong. This distinction is particularly relevant to publicly accountable portfolios where demonstrable service results matter more than technical spectacle.

7 Implementation, research gaps and policy implications

Implementation should begin with portfolio segmentation and a baseline reliability audit. Owners should identify critical cooling, electrical, water and life-safety assets, reconstruct reliable identifiers, reconcile drawings and maintenance records, and examine the distribution of downtime causes. A limited initial deployment should target recurring, costly or high-consequence failures with measurable signals. Success criteria must be agreed before instrumentation: fewer service loss hours, fewer repeat repairs, quicker restoration or a documented reduction in risk exposure. This approach prevents selection of pilot equipment solely because its data are easiest to collect.

A second phase should standardise asset data exchange and connect diagnostics to maintenance execution. Contracts need requirements for point lists, structured fault codes, evidence attachment, recommissioning tests and secure data access. Supervisors should receive training on model uncertainty and escalation, not only software interfaces. Procurement evaluations should ask

whether a solution retains information when the operator or technology provider changes. The ability to export coherent asset history is an operational resilience requirement, not an optional information technology convenience.

Several evidence gaps remain unresolved. First, many fault studies benchmark classification performance but omit repair outcomes, alert workload and consequences of wrong decisions. Second, digital twin investigations often concentrate on technical integration instead of whole-life economic attribution. Third, Saudi-specific longitudinal comparisons across building classes and climates remain limited in the reviewed material. Fourth, cyber risk, vendor dependency and data stewardship are rarely assessed alongside reliability gains. Finally, maintenance studies differ in how downtime and recurrence are defined, obstructing useful comparison across portfolios (Adhikari et al., 2025; Osunsanmi et al., 2025).

Future research should test five propositions: persistent asset identifiers improve end-to-end fault resolution; confidence-calibrated alerts outperform unqualified alarms in technician productivity; criticality-weighted scheduling reduces severe interruptions without increasing total lifecycle cost; verified work order closure lowers repeat failure; and the effect of digital twins is mediated by interoperability and organisational response capacity. Multisite longitudinal designs should publish denominator definitions, negative outcomes and deployment costs. Research should also report whether improvements survive seasonal extremes, model drift and changes in maintenance contractors.

For policymakers and owners, handover requirements could improve the quality of information inherited by operators. Procurement standards can request equipment registers, baseline values and evidence of successful maintenance data transfer. Risk-based oversight should distinguish efficiency analytics from essential life-safety assurance. Portfolio reporting should avoid unsupported claims of avoided failure and instead publish a measurement method. These measures align asset stewardship with broader Saudi digital transformation ambitions while respecting the diversity of building stocks.

8 Limitations

This article is an integrative review, not a registered systematic evidence census or a field evaluation. Thirty dated DOI-identified publications form a purposively bounded corpus, and their heterogeneous designs prevent unbiased quantitative pooling. Some reference metadata are abbreviated where full author lists were not confirmed; exact DOI links are supplied to enable editorial verification. The paper offers no fabricated sample size, search flow totals, maintenance savings or statistically significant effect. Local recommendations are reasoned adaptations that require testing against actual Saudi asset records. Independent plagiarism screening and formal reference manager checks remain necessary before journal submission.

9 Conclusion

Smart building asset management should be judged by dependable service, defensible expenditure and controlled exposure rather than the presence of sensors or digital models. The reviewed literature supports a layered architecture in which credible asset information underpins diagnosis, risk-aware prioritisation, accountable intervention and verification of restored performance. For Saudi Arabia, climate-sensitive cooling systems, mixed building vintages, contractor coordination and procurement data continuity shape practical adoption. The proposed Smart Building Asset Assurance

Framework converts these insights into a staged managerial model with testable performance indicators. Its central claim is conditional: predictive technology improves outcomes only when organisational processes can interpret, act on and learn from the information it provides. Carefully designed longitudinal trials are needed to establish effect size, economic value and transferability across Saudi building portfolios.

References

[1] Adhikari A, Karunaratne T, Sumanarathna N (2025). Machine learning techniques for predictive maintenance of building services: A comprehensive review and research outlook. *Facilities*, 43, 803–817. https://doi.org/10.1108/F-02-2025-0032

[2] AlBalkhy W, Karmaoui D, Ducoulombier L, Lafhaj Z, Linner T (2024). Digital twins in the built environment: Definition, applications, and challenges. *Automation in Construction*, 162, 105368. https://doi.org/10.1016/j.autcon.2024.105368

[3] Alghassab MA (2024). Fuzzy-based smart energy management system for residential buildings in Saudi Arabia: A comparative study. *Energy Reports*, 11, 1212–1224. https://doi.org/10.1016/j.egyr.2023.12.039

[4] Alotaibi S, Alogaili H, Alawwad K, Aljarallah S (2025). The environmental and business benefits of implementing the ISO 50001 energy management system in government buildings: A case study of the Saudi Standards, Metrology and Quality Organization (SASO). *Sustainability*, 17(11), 5131. https://doi.org/10.3390/su17115131

[5] Alrashidi M, Mnasri S, Alqabli M, Alghamdi M, Short M, Williams S, Dawood N, Alkhazi IS, Alrowaily MA (2025). Energy and water management in smart buildings using spiking neural networks: A low-power, event-driven approach for adaptive control and anomaly detection. *Energies*, 18, 5089. https://doi.org/10.3390/en18195089

[6] Asare KAB, Liu R, Anumba CJ, Issa RRA (2024). Real-world prototyping and evaluation of digital twins for predictive facility maintenance. *Journal of Building Engineering*, 97, 110890. https://doi.org/10.1016/j.jobe.2024.110890

[7] Bi J, Wang H, Yan E, Wang C, Yan K, Jiang L, Yang B (2024). AI in HVAC fault detection and diagnosis: A systematic review. *Energy Reviews*, 3(2), 100071. https://doi.org/10.1016/j.enrev.2024.100071

[8] Chen J, Zhang L, Li Y, Shi Y, Gao X, Hu Y (2022). A review of computing-based automated fault detection and diagnosis of heating, ventilation and air conditioning systems. *Renewable and Sustainable Energy Reviews*, 161, 112395. https://doi.org/10.1016/j.rser.2022.112395

[9] Chen Z, et al. (2023). A review of data-driven fault detection and diagnostics for building HVAC systems. *Applied Energy*, 339, 121030. https://doi.org/10.1016/j.apenergy.2023.121030

[10] Dahanayake KC, Sumanarathna N (2021). IoT-BIM-based digital transformation in facilities management: A conceptual model. *Journal of Facilities Management*, 20(3), 437–451. https://doi.org/10.1108/JFM-10-2020-0076

[11] D'Orazio M, Bernardini G, Di Giuseppe E (2023). Predict the priority of end-users' maintenance requests and the required technical staff through LSTM and Bi-LSTM recurrent neural networks. *Facilities*, 41(15/16), 38–51. https://doi.org/10.1108/F-07-2022-0093

[12] Durdyev S, Ashour M, Connelly S, Mahdiyar A (2022). Barriers to the implementation of building information modelling for facility management. *Journal of Building Engineering*, 46, 103736. https://doi.org/10.1016/j.jobe.2021.103736

[13] Ghansah FA (2024). Digital twins for smart building at the facility management stage: A systematic review of enablers, applications and challenges. *Smart and Sustainable Built Environment.* https://doi.org/10.1108/SASBE-10-2023-0298

[14] Gispert DE, Yitmen I, Sadri H, Taheri A (2025). Development of an ontology-based asset information model for predictive maintenance in building facilities. *Smart and Sustainable Built Environment*, 14, 740–757. https://doi.org/10.1108/SASBE-07-2023-0170

[15] Hakimi O, Liu H, Abudayyeh O (2024). Digital twin-enabled smart facility management: A bibliometric review. *Frontiers of Engineering Management*, 11, 32–49. https://doi.org/10.1007/s42524-023-0254-4

[16] Hegazy I (2025). Leveraging advanced digital technologies for smart city development in Saudi Arabia: Opportunities, challenges, and strategic pathways. *International Journal of Low Carbon Technologies*, 20, 834–847. https://doi.org/10.1093/ijlct/ctaf044

[17] Hosamo HH, et al. (2022b). A review of the digital twin technology in the AEC–FM industry. *Advances in Civil Engineering*, 2022, 2185170. https://doi.org/10.1155/2022/2185170

[18] Hosamo HH, Nielsen HK, Alnmr AN, Svennevig PR, Svidt K (2022a). A review of the digital twin technology for fault detection in buildings. *Frontiers in Built Environment*, 8, 1013196. https://doi.org/10.3389/fbuil.2022.1013196

[19] Izaddoost A, Naderpajouh N, Heravi G (2021). Integrating resilience into asset management of infrastructure systems with a focus on building facilities. *Journal of Building Engineering*, 103304. https://doi.org/10.1016/j.jobe.2021.103304

[20] Leygonie R, Motamedi A, Iordanova I (2022). Development of quality improvement procedures and tools for facility management BIM. *Developments in the Built Environment*, 11, 100075. https://doi.org/10.1016/j.dibe.2022.100075

[21] Long W, Bao Z, Chen K, Ng ST, Wuni IY (2024). Developing an integrative framework for digital twin applications in the building construction industry: A systematic literature review. *Advanced Engineering Informatics*, 59, 102346. https://doi.org/10.1016/j.aei.2023.102346

[22] Ma S, Flanigan KA, Bergés M (2024). State-of-the-art review and synthesis: A requirement-based roadmap for standardized predictive maintenance automation using digital twin technologies. *Advanced Engineering Informatics*, 62, 102800. https://doi.org/10.1016/j.aei.2024.102800

[23] Matetić I, Štajduhar I, Wolf I, Ljubic S (2023). A review of data-driven approaches and techniques for fault detection and diagnosis in HVAC systems. *Sensors*, 23, 1. https://doi.org/10.3390/s23010001

[24] Melgaard SP, Andersen KH, Marszal-Pomianowska A, Jensen RL, Heiselberg PK (2022). Fault detection and diagnosis encyclopedia for building systems: A systematic review. *Energies*, 15, 4366. https://doi.org/10.3390/en15124366

[25] Mirnaghi MS, Haghighat F (2020). Fault detection and diagnosis of large-scale HVAC systems in buildings using data-driven methods: A comprehensive review. *Energy and Buildings*, 229, 110492. https://doi.org/10.1016/j.enbuild.2020.110492

[26] Olawale MA, Al Homoud MS, Abdou AA, Mohammed MA (2025). A comprehensive assessment of the impact of the Saudi Building Code on energy performance of residential buildings in Saudi Arabia. *Energy and Buildings*, 349, 116524. https://doi.org/10.1016/j.enbuild.2025.116524

[27] Osunsanmi TO, Okafor CC, Aigbavboa CO (2025). Critical success factors for implementing smart maintenance in the fourth industrial revolution era: A bibliometric analysis within the built environment. *Journal of Facilities Management*, 23, 209–230. https://doi.org/10.1108/JFM-01-2023-0006

[28] Valinejadshoubi M, Moselhi O, Bagchi A (2022). Integrating BIM into sensor-based facilities management operations. *Journal of Facilities Management*, 20, 385–400. https://doi.org/10.1108/JFM-08-2020-0055

[29] Zhang A, Yang J, Wang F (2023). Application and enabling digital twin technologies in the operation and maintenance stage of the AEC industry: A literature review. *Journal of Building Engineering*, 80, 107859. https://doi.org/10.1016/j.jobe.2023.107859

[30] Zhang L, Leach M, Bae Y, Cui B, Bhattacharya S, Lee S, Im P, Adetola V, Vrabie D, Kuruganti T (2021). Sensor impact evaluation and verification for fault detection and diagnostics in building energy systems: A review. *Advances in Applied Energy*, 3, 100055. https://doi.org/10.1016/j.adapen.2021.100055