

Digital Readiness Tracking for Operational Readiness in Mission-Critical Infrastructure: A Practitioner Framework for Near-Real-Time, System-Based Startup Verification

Rushi Jariwala

RUS Industrial, LLC, USA

Abstract: Mechanical completion certifies that a facility has been built to specification. It does not certify that the facility, its operating organization, and its documentation are ready to be started up and operated safely. On large industrial and mission-critical projects, the space between those two states tends to be where schedule overruns, disputed handovers, and early-life failures concentrate. This article offers a practitioner framework for operational readiness (OR) that treats readiness as a planning-stage objective spanning equipment, documentation, training, spare parts, maintenance procedures, and safety systems, and it sets out what a digital tracking system would need to capture and deliver to make readiness verifiable in near real time. The case evidence presented is drawn from turnover and commissioning readiness records quality, testing, equipment, and turnover-documentation streams the author directly verified and is narrower than the full OR concept the framework describes; it does not extend to verified training, spare-parts, or maintenance-procedure records. Four elements organize the framework early operations involvement, integrated readiness planning, system-based readiness metrics, and system-based completion tracking around one principle: data should be structured by commissioning system rather than by construction discipline. For each element, the article describes the data fields, aggregation rules, and reporting cadence a dashboard or digital twin would require. The framework is illustrated against four capital projects executed by the author, among them a 121-system petrochemical expansion (765 hydrotest, 824 equipment, and 5,264 instrumentation packages), a methanol facility of roughly 25,000 welds, and a 78-module, \$120M data center program. Results are reported qualitatively, in terms of visibility, accountability, and startup predictability. The analysis points to a gap between current digital-tracking research, which emphasizes sensing and progress monitoring, and the system-level data organization practitioners appear to need to verify readiness across quality, testing, equipment, and commissioning at once.

Keywords: Operational readiness; commissioning and startup; system-based completion tracking; turnover packages; digital twin; readiness dashboard

1. INTRODUCTION

Large capital projects in the process and mission-critical sectors answer two very different questions. One asks whether the facility has been constructed in accordance with its drawings, codes, and specifications. The other asks whether the facility can be started up and operated as an integrated system without interruption. Construction management practice has tended to concentrate on the first question, expressed through mechanical completion and punch list closure, and to treat the second as a downstream consequence that should follow once construction finishes. Field experience across two decades suggests that assumption does not hold reliably.

Operational readiness (OR) may be better understood as a strategy distinct from, and broader than, mechanical completion. Mechanical completion is a construction-end certification that physical work conforms to design. OR is a planning-stage objective, and it asks whether equipment, documentation, spare parts, trained operators, maintenance procedures, and safety systems will all be available and verified at the moment a commissioning system is energized.



That difference matters because the two states are certified against different data. Verification of mechanical completion runs discipline by discipline piping, electrical, instrumentation while startup happens system by system, and a system is not started until every discipline inside its boundary is complete and turned over together.

Readiness at the construction stage has been studied directly. The Construction Industry Institute compiled 228 readiness factors across 15 categories, derived from 80 projects (41 construction-ready versus 39 not-ready), which established that readiness is measurable and that unreadiness is common [1]. Building on that, a comprehensive model for construction readiness assessment has been proposed to score projects before authorization [2]. Ignoring readiness carries documented costs: premature construction starts have been shown to generate downstream disputes and litigation whose root cause traces to work begun before its predecessors were finished [3]. Similar logic seems to apply at the other end of the project, where premature startup declaring a system ready before its readiness data is closed transfers risk into the operating phase.

Mission-critical infrastructure makes that risk concrete. Data center operators report that a majority of significant outages carry direct costs above \$100,000, and that a substantial share stem from human error traceable to procedures that were not followed [4]. Readiness failures, on this evidence, are not confined to the schedule; they can resurface as early-life operational failures once a facility is live.

The mission-critical setting adds a sharper edge to the problem. A delayed system startup in a petrochemical plant is costly but usually recoverable inside the commissioning window. A data center carrying live compute load has far less tolerance for an unverified system reaching service, since the failure mode is not a late startup but an outage against an operating customer. The readiness discipline that has always mattered in process construction appears to carry a steeper penalty for error in the mission-critical sector, which is part of why system-based verification, over discipline-based progress reporting, looks worth formalizing.

What this article sets out to do is specify, from direct field experience, what a digital tracking system would need to deliver for OR verification. No new sensing technology or digital twin architecture is proposed here. The problem is approached from the practitioner side: given how readiness is actually verified on a turnover-driven project, what data structure, aggregation logic, and reporting cadence would a dashboard or digital twin have to provide to make readiness visible in near real time? The framework that follows is built around that question and illustrated against four projects the author has executed.

2. RELATED WORK

The case for digitalizing construction data rests partly on the sector's productivity record. Construction labor-productivity growth has averaged only 1% a year over two decades, versus 2.8% for the total world economy, a shortfall often attributed to fragmented information and manual coordination [5]. Readiness verification depends on assembling quality, testing, and equipment records across many parties, which makes it exactly the kind of information-intensive task digitalization is expected to help.

Digital twins for construction and facilities management have drawn sustained attention. Researchers have mapped directions for a semantic construction digital twin [6] and traced the evolution from building information models (BIM) toward digital twins as intelligent building representations [7]. Related efforts have built frameworks and common data environments for BIM-based facilities management using open standards [8], and have proposed typologies of the information handover requirements that link construction to operations [9]. Sensor-level integration has moved forward too, including methods to enrich BIM semantically with sensor descriptions for built-environment monitoring [10]. This body of work shows that a facility's static and dynamic data can be represented digitally, though its orientation leans toward the operations phase and toward interoperability more than toward verifying readiness during construction turnover.

Real-time tracking of site work has grown along several lines. Indoor resource positioning has been used to infer task progress from the movement of labor and materials [11], and BIM-IoT integration has been applied to real-time quality monitoring on site [12]. Wider frameworks for real-time tracking and analysis have been proposed to

consolidate these signals [13]. Computer-vision methods now support progress monitoring from site imagery [14] and worker tracking through deep-learning instance segmentation in off-site construction [22]. Such approaches are strong at answering how much physical work has been done, yet they observe work discipline by discipline and location by location, not commissioning system by commissioning system.

Predictive analytics has been turned toward schedule and startup performance as well. Machine learning paired with the Last Planner System has been used to predict schedule performance [15], and data-driven methods have folded field submittals into project scheduling [16]. Commissioning and startup have their own dedicated study from the Construction Industry Institute, which identified 16 critical success factors across 45 indicators, validated against 26 capital projects [17]. That research confirms startup success is predictable from front-end factors, but it stops short of prescribing the running data structure needed to track those factors during execution.

Modular and offsite construction, now central to data center delivery, has generated its own quality and readiness work. Digital twin methods have been proposed for production estimation, scheduling, and real-time monitoring in offsite construction [18], and quality-information management systems have been designed for modular construction factories [19]. Information loss at handover is a persistent concern examined from the BIM side, both in developing models for building handover, operation, and maintenance [20] and in frameworks aimed specifically at managing handover information loss [21].

One observation cuts across these strands. Each handles a single slice of the readiness picture BIM interoperability, physical progress, schedule prediction, factory quality, or handover completeness yet the slices are mostly studied apart. Verifying operational readiness appears to require all of them together, resolved to the same commissioning system: the quality record (welds, NDE, MTR), the testing record (hydrotest and pneumatic packages), the equipment record (static and rotating turnover), and the commissioning record (loop checks, energization prerequisites). No common structure shows up in the literature into which those four records collapse to a single per-system verdict, which is the artifact a startup manager tends to reach for.

This article is related to, and should be read alongside, the author's companion manuscript on an Integrated Turnover Management System, which is drawn from the same four projects and shares their underlying scale figures [23]. The two articles address different problems and should not be treated as reporting the same contribution twice. The companion article proposes a readiness-measurement model: a scored, gated method (including a system-readiness-percentage calculation, per-element acceptance criteria, sequential status gates, and a supporting KPI set) for determining, at a point in time, how ready a given commissioning system is. This article instead specifies the digital data architecture a tracking system would need to hold and refresh to make that kind of readiness verdict computable in the first place: the data dictionary, aggregation rules, dependency logic, and reporting cadence a dashboard or digital twin would require. Put differently, the companion article answers "how ready is this system, and how do we score that," while this article answers "what would the underlying data system have to look like for that score to be computed from live records rather than assembled by hand." Readers interested in the scoring method itself, rather than the data architecture beneath it, should consult the companion article directly.

Read together, this literature converges on sensing, interoperability, and progress monitoring. Less well covered is a specification of the data structure and cadence a practitioner needs to verify system-level operational readiness across quality, testing, equipment, and commissioning at the same time. Digital tools can report that a weld was made, that a worker was present, or that a BIM object exists; they do not, by default, report that commissioning system 42 has every hydrotest package closed, every equipment turnover accepted, every instrument loop checked, and every open NCR resolved. The distance between discipline-organized sensing and system-organized readiness is the gap this article works on.

Table 1. Literature Streams Informing the Digital Readiness Tracking Framework

Literature Stream	Relevance to Framework
-------------------	------------------------

Digital twins for construction/FM	Digital representation of facility state; leans toward operations/interoperability, not turnover-stage readiness
Real-time site tracking (IoT/CV)	Observes physical progress by discipline/location, not by commissioning system
Predictive analytics for schedule/startup	Predicts startup success from front-end factors; does not prescribe a running data structure
Modular/offsite quality and handover	Factory quality control and BIM-to-ops handover-loss risk, directly relevant to modular data center delivery
Mission-critical reliability stakes	Quantifies the cost of readiness failures once a facility is operating

3. OPERATIONAL READINESS FRAMEWORK AND DIGITAL TRACKING REQUIREMENTS

Four elements make up the framework: early operations involvement, integrated readiness planning, system-based readiness metrics, and system-based completion tracking. Each is described first as a practice and then as a set of requirements a digital tracking or dashboard system would have to satisfy. The quantitative examples that follow are drawn from the author's direct project execution experience, and they illustrate scale and data structure, not experimental results.

3.1 *Early operations involvement*

Readiness starts in planning, well before construction ends. When operations, commissioning, and maintenance personnel are engaged during design and early execution, the definition of "complete" gets negotiated up front: turnover package contents, system boundaries, and acceptance criteria are agreed before the first weld map is issued. Defer that engagement, and construction tends to optimize for discipline completion, leaving the operating organization with a facility that is built but not verifiably ready.

For a digital tracking system, early involvement translates into a fairly specific requirement. The data model should encode commissioning system boundaries as first-class objects at project setup, before construction data begins to accumulate. Each physical component line, weld, instrument, equipment tag would carry a commissioning system attribute from the moment it enters the model. Absent that attribute, discipline data can be aggregated to the system level only through manual reconciliation, which rarely keeps pace with the field. The framework asks that system boundaries be defined and loaded as a controlling structure, and that operations stakeholders see that structure from the outset.

System boundaries are also where the operations and construction organizations most often disagree, one more argument for fixing them early. A boundary drawn for construction convenience along a pipe rack, say, or a contract package seldom lines up with the boundary along which a system gets commissioned and operated. Reconcile the two only at turnover, and the reconciliation itself turns into a source of delay and dispute, since records filed against the construction boundary have to be re-sorted against the commissioning boundary by hand. Set the commissioning system boundary at project setup, hold both organizations to it, and that late reconciliation comes off the critical path, leaving the digital tracking system a stable key against which every record can be filed from the day it is created. In the author's experience, turnover ran most predictably on projects where this boundary discussion happened during design reviews, well before construction mobilized.

3.2 *Integrated readiness planning*

Readiness is less a single checklist than the simultaneous satisfaction of several parallel streams: equipment installation and verification, documentation and turnover packages, operator training, spare parts, maintenance procedures, and safety systems. A facility is ready to start a given system only when all streams are complete inside

that system's boundary. These streams are usually owned by different parties and tracked in different tools, which is a large part of why readiness gets asserted more often than it gets demonstrated.

Real project structure conveys the scale better than description does. On the INEOS Styrolution 100KT ASA expansion, execution was organized into 121 commissioning systems, against which the project managed 765 hydrotest packages, 824 equipment packages, and 5,264 instrumentation packages (based on the author's direct project execution experience). Each of those thousands of interrelated package records belongs to one commissioning system, and each carries its own status across the readiness streams. A digital tracking system built for integrated readiness planning would therefore hold, for every commissioning system, the roll-up status of each parallel stream, and would express readiness as the logical conjunction of those streams rather than as an average or a single percent-complete figure. A system reported at 95% on piping but missing its rotating-equipment MTR review is not 95% ready; it is not ready.

That requirement carries into the linkage between records. A hydrotest package should not close until its constituent welds have passed non-destructive examination and their MTR review is accepted; an instrument loop should not be checked until its associated equipment is set and its documentation turned over. The framework asks a digital system to encode these dependencies explicitly, so a stream cannot report readiness while an upstream record it depends on stays open. This mirrors, at the readiness end, the predecessor logic whose violation drives premature-start disputes [3].

Non-physical streams deserve equal standing in the data model. Operator training, spare-parts receipt and staging, approved maintenance procedures, and verified safety systems are readiness prerequisites that no amount of physical completion can stand in for, and they are often tracked outside the construction system altogether in training rosters, warehouse logs, and procedure libraries that never resolve to a commissioning system. A tracking system built for operational readiness would want to fold these streams into the same system-based structure, so a system's readiness verdict reflects whether its operators are trained and its procedures approved, not just whether its pipe has been tested. The human-error findings in mission-critical operations [4] make that look less like a peripheral requirement than a central one.

3.3 System-based readiness metrics

At the center of the framework sits a claim with real consequences: readiness should be measured by the commissioning system, not by construction discipline. Discipline metrics total welds complete, total instruments installed answer a construction question, and they are the natural output of most tracking and sensing tools, the computer-vision and IoT approaches surveyed earlier among them [11,12,14]. They struggle with the readiness question, because a facility does not start up all its piping at once. It starts up one system at a time, and a system's readiness depends on the completion of every discipline inside its boundary, not on the project-wide completion of any single discipline.

Consider what that distinction does to the data structure. A project may sit at 90% complete on welds and still have zero systems ready if the remaining 10% of welds fall across the critical-path startup systems. What a startup manager needs is not the discipline percentage but the ordered list of systems, each with a conjunctive readiness status and a resolved list of what remains inside its boundary. A digital tracking system would compute readiness bottom-up, then: from component records, to package status, to stream roll-up, to a per-system readiness verdict, presented in startup sequence.

Cadence belongs in the requirement too. Readiness on an active turnover front shifts daily as packages close and NCRs are raised, so a dashboard that refreshes monthly cannot really govern a startup sequence measured in weeks. The framework asks that per-system readiness be recomputed continuously from the underlying records, so the readiness dashboard reflects field state close to real time, not at reporting-period boundaries. Here digital tracking seems to offer its clearest advantage over manual turnover logs, and much of the surveyed sensing literature looks well placed to supply data provided that data is aggregated to the system rather than the discipline.

3.4 System-based completion tracking

The last element puts the metric into practice: completion is tracked, reported, and certified at the system level through turnover packages assembled by system boundary. A turnover package is the evidentiary record that a system is ready with its hydrotest and pneumatic packages, weld maps, MTR reviews, ITP/QCP sign-offs, punch list closure status, and open NCR/CAR log, all bound to the commissioning system. Workface sequencing should be planned to deliver complete systems in startup order, not to maximize discipline throughput.

For a digital system, the requirement is that the turnover package behave as a live, queryable object rather than a static document assembled at the end. Each element of the package would link to its underlying records, so its status updates as field work closes, and so punch list closure and NCR resolution stay visible against the system boundary in near real time. A dashboard should draw a clear line between construction-complete and readiness-complete states, since the difference between them is more or less the risk operational readiness exists to manage.

Two reporting behaviors seem to follow from this structure. The dashboard might favor exception-based reporting: instead of leading with completion percentages, it would surface, for each system in startup sequence, the specific open items blocking readiness an unresolved NCR, an outstanding MTR review, an incomplete loop check. Performance Factor (PF) tracking would ideally hang off the same system structure, so the earned progress reported to project controls and the readiness reported to the startup organization come from one dataset rather than two reconciled later. A divergence between them high earned progress against low readiness on the same system is arguably the signal most worth watching.

Pulling these threads together, the framework's data requirements reduce to five capabilities a digital tracking or digital twin system would need to provide: (1) commissioning system boundaries as first-class objects loaded at setup; (2) a system attribute on every component and package record; (3) parallel readiness streams rolled up per system and combined by logical conjunction; (4) explicit dependency logic that keeps a stream from reporting readiness over an open predecessor; and (5) continuous, per-system readiness computation presented in startup sequence. None of these calls for a particular sensor or algorithm. Each is a statement about how the data has to be organized for readiness to be verifiable.

Table 2. Five Data Requirements for a Digital Operational Readiness Tracking System

#	Requirement	Framework Element
1	Commissioning system boundaries encoded as first-class objects at project setup	Early operations involvement
2	A commissioning-system attribute on every component and package record	Integrated readiness planning
3	Parallel readiness streams rolled up per system and combined by logical conjunction, not averaged	System-based readiness metrics
4	Explicit dependency logic preventing a stream from reporting readiness over an open predecessor	System-based readiness metrics
5	Continuous, per-system readiness computation presented in startup sequence	System-based completion tracking

3.5 Data architecture, dashboard structure, and a worked example

The requirements in Table 2 imply a specific data architecture, not just a design principle, and this subsection makes that architecture concrete. Four elements complete the specification: a record-level data dictionary, a status-and-dependency logic for readiness computation, a dashboard structure that presents the result, and an update cadence that keeps it current. No software was built to implement this architecture; what follows is the schema the framework requires, illustrated against a single representative system from the case projects.

3.5.1 Data dictionary

Every record in the system resolves to a commissioning system through a shared `system_id` key, carries a stream tag identifying which readiness stream it belongs to, and carries an explicit list of dependency record_ids it cannot report ready ahead of. Table 2b sets out the minimum field set.

Table 2b. Minimum Data Dictionary for a System-Based Readiness Record

Field	Definition	Example value
<code>system_id</code>	Unique identifier for the commissioning system; the primary key every other record resolves to	SYS-042
<code>system_name</code>	Human-readable system name used in startup sequencing	Reactor Feed System 42
<code>startup_sequence_no</code>	Integer rank giving the system's position in the planned startup order	17
<code>record_type</code>	Package category: hydrotest, pneumatic, equipment, instrumentation, weld, loop-check, NCR, MTR	instrumentation
<code>record_id</code>	Unique identifier for the individual package or record	INST-3381
<code>stream</code>	Which of the parallel readiness streams the record belongs to (equipment, documentation, training, spare parts, maintenance, safety)	equipment
<code>dependency_ids</code>	List of record_ids that must close before this record can report ready (e.g., NDE and MTR prerequisites for a hydrotest package)	[WELD-1187, MTR-0042]
<code>status</code>	Record-level state: open, in-progress, closed-pending-review, closed	closed
<code>last_updated</code>	Timestamp of the most recent status change, used to drive cadence	2026-06-14T08:12Z
<code>authority</code>	Role or party accepting the record (e.g., QC inspector, commissioning engineer)	Commissioning Engineer

3.5.2 Readiness-status logic

A system's readiness status is computed bottom-up and by logical conjunction, not by averaging. At the record level, a record is ready only if its own status is closed and every record in its `dependency_ids` list is also closed; a hydrotest package cannot report ready while an upstream weld or MTR record it depends on remains open. At the stream level, a stream (equipment, documentation, training, spare parts, maintenance, safety) is ready only if every record tagged with that stream and that `system_id` is ready. At the system level, the system is ready only if every stream is ready: `ready(system) = ready(equipment) AND ready(documentation) AND ready(training) AND ready(spare_parts) AND ready(maintenance) AND ready(safety)`. A single open record in any stream holds the whole system at not-ready, which is the conjunctive logic Section 3.3 argues for and which an averaged percent-complete figure cannot express.

3.5.3 Dashboard structure and update cadence

The dashboard the architecture supports is organized in startup sequence, not alphabetically or by discipline: each row is one commissioning system, ordered by `startup_sequence_no`, with a single ready/not-ready verdict per system and, for any not-ready system, the specific open records blocking it rather than a completion percentage. Because every record carries a `last_updated` timestamp, the readiness verdict is recomputed whenever an underlying record changes state, rather than on a fixed reporting cycle; in practice this means the dashboard reflects field state within the latency of the source systems feeding it (e.g., the turnover-tracking tool or quality database), which the case projects' manual, spreadsheet-based processes could not sustain below a multi-day cycle. This is the near-real-time property the framework targets: continuous recomputation from source records, not a claim that any sensor network or live operational deployment was built or tested.

3.5.4 Worked example

Consider Reactor Feed System 42 on a project structured like the INEOS Styrolution expansion (Section 5.1). The system carries an equipment-turnover record (closed), a hydrotest package with two dependency records for weld NDE and MTR review (both closed, so the hydrotest package reports closed), and an instrumentation loop-check record that depends on the equipment-turnover record and remains open pending a vendor calibration certificate. Applying the logic in Section 3.5.2, the equipment and documentation streams read ready, but the instrumentation stream does not, because one dependent record is open; the system-level verdict is therefore not-ready, and the dashboard surfaces the specific open item (the pending calibration certificate) rather than a system completion percentage such as "92% complete," which would obscure the fact that the system cannot yet be started. This is the mechanism by which the architecture converts the record-level data in Table 2b into the per-system, startup-sequenced verdict the case narratives in Section 5 describe qualitatively.

4. METHODOLOGY

This article follows a retrospective multi-project practitioner case study design. The framework in Section 3 was developed inductively from the author's direct execution record and is illustrated, not statistically tested, against four capital projects. No experimental manipulation, control group, or randomized comparison is involved; the projects were not selected or instrumented for research purposes at the time they were executed. This section sets out project selection, the author's role on each project, the records reviewed, the steps taken to verify the figures reported, the confidentiality controls applied, and the limitations that follow from this evidence base.

4.1 Project selection

The four projects discussed in Section 5 (INEOS Styrolution 100KT ASA expansion, Braskem polypropylene recovery program, NatGasoline methanol facility, and the Core Scientific / RUS Industrial modular data center program) were selected because they are the capital projects on which the author held a quality- or project-management role with direct, first-hand visibility into turnover and readiness data, and because together they span the sector, delivery-method, and scale variation the framework claims to generalize across: conventional stick-built petrochemical construction, a brownfield recovery program operating alongside a live plant, a large greenfield methanol facility, and modular data center delivery. No project on which the author lacked direct execution

responsibility was considered for inclusion, and no project was excluded from the four on evidentiary or outcome grounds; the selection is exhaustive of the author's qualifying project record at the time of writing, not a sample drawn from a larger pool.

4.2 Author's role on each project

The author's role varied by project and is disclosed here because it bears on what each project's figures represent. On the INEOS Styrolution 100KT ASA expansion the author held a site quality control and commissioning-support role with direct responsibility for turnover package status across the 121 commissioning systems discussed in Section 5.1. On the Braskem polypropylene recovery program the author held a quality/turnover role scoped to the recovery systems described in Section 5.2, with visibility into the hydrostatic/pneumatic test and equipment turnover packages reported there. On the NatGasoline methanol facility the author held a quality and project-controls role with visibility into weld, hydrotest, and equipment turnover status across the project reported in Section 5.3. On the Core Scientific / RUS Industrial modular data center program the author holds a quality/project-management role, and the figures in Section 5.4 reflect that ongoing engagement as of the data cutoff stated there. In each case the figures reported are those the author's role gave direct access to verify; figures outside that scope are not reported.

4.3 Records reviewed and data verification

The quantitative figures reported in Section 5 and Table 3 (package counts, weld counts, system counts, module counts, manpower, and manhours) are drawn from the author's direct project execution experience and from project turnover, quality, and progress records the author had contemporaneous access to in each of the four roles above. Where a figure could be cross-checked against a project turnover log, punch list system, or progress-reporting record available to the author at the time of writing, it was; where a figure could not be independently cross-checked against a surviving record, or where its disclosure was not authorized by the employer or client of record, it is marked in Table 3 as "not separately disclosed" or reported with an authorization caveat, and no figure is presented as more precise than the underlying record supports. Approximate figures are reported as approximate ("approx.", "roughly") rather than as exact counts. No figure in this article was estimated, extrapolated, or reconstructed from memory alone where a contemporaneous record was unavailable; in that situation the figure is omitted or qualified rather than supplied.

4.4 Confidentiality and disclosure controls

Project, employer, and client names are used where their disclosure is a matter of public record (e.g., published project announcements) or where prior disclosure by the employer or client permits it; commercial figures such as contract value, and workforce figures such as peak manpower and manhours, are disclosed only to the extent the author understands the employer or client to authorize for public disclosure, and are otherwise reported qualitatively or omitted. No proprietary drawings, specifications, financial statements, or personally identifying information about other individuals are reproduced or described in this article. Where a quantity could reveal commercially sensitive information beyond what is already authorized for disclosure, the article reports it in aggregate or qualitative terms instead.

4.5 Limitations of the evidence

This is a single practitioner's retrospective case record, not a controlled study, and it carries the limitations that follow from that design. There is no comparison group of projects run on discipline-based tracking against which the framework's qualitative benefits (visibility, accountability, startup predictability) could be measured; the article does not claim to isolate the framework's causal effect on any outcome, and the quantitative figures are used to convey project scale and data structure, not to demonstrate performance improvement. Because the author executed all four projects, the account is a single-observer record and is subject to the hindsight and selection effects any practitioner's retrospective account carries; a different practitioner on the same projects might weight the same events differently. The evidence is also confined to U.S. industrial and mission-critical contexts and to the delivery models (stick-built petrochemical, brownfield recovery, greenfield methanol, modular data center) the four projects represent; it does not

speak to other regulatory environments or delivery models. These limitations, and their implications for future validation work, are returned to in Section 7.

5. CASE APPLICATION

The framework is illustrated below against four projects the author has executed. They differ in sector, delivery method, and scale, and between them they span conventional stick-built petrochemical construction, a brownfield recovery program, a large greenfield methanol facility, and modular data center delivery. Results are described qualitatively in visibility, accountability, startup predictability rather than through invented performance percentages.

Table 3. Cross-Project Scale (Author's Direct Project Execution Data)

Project	Sector	Approx. Value	Key Records	Commissioning Structure
INEOS Styrolution 100KT ASA Expansion	Petrochemical	Not separately disclosed	765 hydrotest packages; 824 equipment packages; 5,264 instrumentation packages	121 commissioning systems
Braskem Polypropylene recovery program	Petrochemical (recovery)	Not separately disclosed	Approx. 725 hydrostatic/pneumatic test packages; approx. 120 equipment turnover packages	Brownfield recovery systems with live-plant interfaces
NatGasoline Methanol Facility	Methanol production	Not separately disclosed	Approx. 25,000 welds; 830 hydrotest packages; 170 static + 73 rotating equipment units	Not separately itemized
Core Scientific / RUS Industrial modular data center program	Mission-critical data center	~\$120 million	78 modules; ~440 peak manpower; 600,000+ manhours (as of July 2026 data cutoff, to the extent authorized for disclosure)	78-module structure, factory + site

Note: figures are the author's direct project execution records; entries marked "not separately disclosed/itemized" reflect the scope the author tracked and can verify for that project.

5.1 INEOS Styrolution 100KT ASA expansion

As noted, the ASA expansion was organized into 121 commissioning systems carrying 765 hydrotest packages, 824 equipment packages, and 5,264 instrumentation packages (based on the author's direct project execution experience). A record count of this size resists governance by discipline totals: the sheer volume of instrumentation packages means a small residual percentage can represent hundreds of open records, and their distribution across systems determines what can actually be started.

Applying the framework, completion was tracked against the 121-system structure, with hydrotest, equipment, and instrumentation records each carrying their system attribute so readiness could be rolled up per system. The practical payoff showed up in visibility and accountability. Because every open package resolved to a named system, the residual work in front of any given startup system was explicit, not buried in a discipline average. Turnover packages assembled by system boundary made the difference between construction-complete and readiness-complete

legible to the commissioning organization, which appears to be a condition under which startup predictability improves. The instrumentation count, in particular, illustrates why manual, discipline-organized tracking becomes hard to sustain and why a system-attributed digital structure earns its keep.

Workforce sequencing on the expansion was planned so crews closed out systems in an order aligned with the commissioning sequence, which trimmed the population of near-complete-but-not-ready systems waiting on scattered residual work. Where residual instrumentation packages clustered on a critical system, the system-attributed structure tended to expose that cluster early enough to redirect effort, well ahead of turnover.

5.2 Braskem polypropylene recovery program

A different challenge defined the Braskem recovery program: a brownfield environment with live-plant interfaces, tie-ins, and operating systems, where sequencing and reconciliation of readiness against operating constraints, rather than raw record volume, tended to dominate. The recovery scope carried approximately 725 hydrostatic and pneumatic test packages and approximately 120 equipment turnover packages (based on the author's direct project execution experience). Recovery work of this kind turns on bringing specific systems back to a ready state while the surrounding facility keeps operating, which puts a premium on tight system boundaries and unambiguous readiness verdicts.

The framework's contribution here sat in accountability and boundary discipline. Defining what fell inside each recovery system's boundary, and holding readiness to the conjunction of its streams rather than to scattered discipline progress, helped keep recovery scope from bleeding into adjacent operating systems. A digital tracking system in this setting would want the same system-first data model, with particular weight on the dependency logic that keeps a system from being declared ready while an open NCR or an incomplete turnover element remains inside its boundary.

With the plant partly in service, each recovery system's readiness verdict also had to account for interfaces with live systems isolation points, tie-ins, and the safety systems needing verification before energization. The dependency logic extended fairly naturally to those interfaces: a recovery system would not read as ready while an interface to an operating system stayed unverified.

5.3 NatGasoline methanol facility

Scale returned to the foreground at the NatGasoline methanol facility, a large greenfield project of roughly 25,000 welds, 830 hydrotest packages, and 170 static plus 73 rotating equipment units (based on the author's direct project execution experience). At this weld count the gap between discipline completion and system readiness is about as visible as it gets: a facility can read as nearly complete on welding while the specific welds remaining sit on critical startup systems, and the rotating equipment few in number but decisive for startup can gate readiness regardless of overall piping progress.

This is where the framework's insistence on bottom-up, per-system readiness computation looks most justified. Weld completion meant something only once resolved to hydrotest packages and then to commissioning systems; the 73 rotating equipment units needed their own readiness verification alignment, MTR review, turnover before their host systems could count as ready, independent of the surrounding weld percentage. The case suggests a digital tracking system should weight readiness by system criticality and equipment dependency, not by aggregate count, and should treat MTR review and equipment turnover as first-class elements of the system readiness verdict.

How readiness was sequenced turned partly on the static-versus-rotating split. Static equipment vessels, exchangers, tanks was generally verified through its associated piping and hydrotest packages. The 73 rotating units carried extra prerequisites: lube systems, alignment records, and vendor sign-offs to be closed before the driven system could run. Treating those prerequisites as first-class readiness elements, not as equipment-discipline line items, kept rotating-equipment readiness from hiding behind high piping completion.

5.4 Core Scientific modular data center program

This program reflects an employer contract scope of approximately \$120M delivering 78 OptiStax modules, at peak manpower of about 440 and more than 600,000 manhours (based on the author's direct project execution experience, to the extent authorized for public disclosure, as of a July 2026 data cutoff). Modular data center delivery departs from process construction in that much of the work is completed offsite in a controlled factory and then integrated on site, which shifts part of the readiness problem onto factory quality control and onto the integrity of information handed from factory to field the very handover-loss problem examined in the modular and BIM literature [18,19,20,21].

Here the framework maps fairly naturally onto a module-and-system structure. Each module carries its own quality and turnover records, and site readiness depends on both the factory-verified state of each module and the integration work that ties modules into operable systems. Mission-critical reliability raises the stakes of getting this right, since readiness failures in the sector can resurface as operational outages with direct financial consequences [4]. The requirement that turnover packages be live, system-attributed objects speaks to the handover-loss risk fairly directly: link factory quality records into the same system-based structure that governs site readiness, and the information does not have to be rebuilt at the factory-to-field boundary. This program is probably where the case for a digital twin-integrated readiness dashboard, as against document-based turnover, is strongest.

Roughly 440 peak manpower and more than 600,000 manhours give a sense of the coordination load a system-based structure has to absorb without collapsing back into discipline reporting. Modular delivery splits the readiness question across two places at once: the factory, where module quality is verified before shipment, and the site, where modules are set, interconnected, and commissioned into operable systems. A shared system-based schema spanning both would let a module accepted at the factory arrive with its quality record already attached to the site system it serves, instead of as a separate document set to be re-verified on receipt.

Read across the four projects, a common pattern shows through despite the differences in sector and delivery. The quantity and diversity of records in each case more than 5,000 instrumentation packages on one, roughly twenty-five thousand welds on another, seventy-eight factory-built modules on a third outran what discipline-organized reporting could turn into a readiness answer. What constrained startup, project after project, was not the project-wide completion percentage of any single discipline but the residual work concentrated on specific systems, and that residual became actionable only once it resolved to a named commissioning system. Which stream tended to gate readiness varied from job to job: instrumentation on the ASA expansion, interface verification on the Braskem recovery, rotating equipment on NatGasoline, factory-to-field handover on the Core Scientific program. A tracking system general enough to serve all four would have to bend to whichever stream dominates on a given project, which seems to argue for a uniform system-based structure over a discipline-specific tool tuned to one sector.

6. DISCUSSION: LESSONS LEARNED

A few lessons recur across these projects about what digital readiness tracking has to get right. The first, and the most important, is data organization. A tracking system that stores completion and quality data by discipline can report construction progress but struggles to verify operational readiness, because readiness is a system-level conjunction the discipline structure does not express. The single most consequential design decision seems to be making the commissioning system the organizing key of the data model a decision best taken at project setup, since retrofitting it means reconciling every record by hand. That reading lines up with construction-readiness research showing readiness is measurable and predictable when the right factors are tracked from the front end [1,2,17].

Real-time visibility is the second lesson. A readiness dashboard earns its value by reflecting field state as it changes, so the startup sequence gets governed by current data over month-old reports. The surveyed sensing, IoT, and computer-vision literature looks well positioned to feed such a dashboard [10,11,12,14,22], though only if its outputs are aggregated to the commissioning system; discipline-level or location-level sensing, however accurate, does not on its own answer the readiness question. What a practitioner framework may add to that literature is the specification of the aggregation target.

Leadership ownership is the third lesson, and it is not really a technical one. A readiness dashboard changes behavior only when the operating and construction organizations are jointly accountable to it, which in turn asks for early operations involvement and a shared definition of readiness agreed before construction data piles up. The human-error findings in mission-critical operations [4] are a reminder that verified procedures and trained personnel belong to readiness, not alongside it; a dashboard that tracks physical completion while ignoring documentation, training, and procedures is measuring the wrong conjunction [24, 25].

A fourth lesson concerns the limits of automation. Sensing and computer-vision tools can observe physical state, and predictive models can forecast schedule outcomes, but neither seems to substitute for the human judgment that closes a turnover package accepting an MTR review, dispositioning an NCR, signing an ITP hold point. A digital readiness dashboard is probably better cast as organizing the evidence behind that judgment, and making the residual clear, than as replacing the judgment itself. On the projects reviewed here, the value of system-based tracking lay less in automating decisions and more in getting the right decisions in front of the right people while there was still time to act on them [26].

Table 4. Lessons Learned for Digital Readiness Tracking

Lesson	Practical Implication
Data organization is the single most consequential design decision	Make the commissioning system the organizing key at project setup, not a retrofit
Real-time visibility is what a dashboard is for	Aggregate sensing/IoT/CV outputs to the commissioning system, not discipline or location
Leadership ownership determines whether the dashboard changes behavior	Joint accountability between operations and construction, agreed before data accumulates
Automation organizes evidence; it does not replace judgment	Design dashboards to surface the residual clearly, not to auto-decide turnover acceptance

7. CONCLUSION AND FUTURE WORK

This article has argued that operational readiness is better treated as a planning-stage objective distinct from mechanical completion, and it has set out, from direct field experience, what a digital tracking system would need to deliver to make readiness verifiable in near real time. The framework's four elements early operations involvement, integrated readiness planning, system-based readiness metrics, and system-based completion tracking come down to one design requirement: readiness data should be organized by commissioning system, rolled up across parallel streams by logical conjunction, and computed continuously in startup sequence. Illustrated against four capital projects across petrochemical, methanol, and modular data center construction, the framework separates construction-complete from readiness-complete and makes the residual work in front of each startup system explicit.

Limitations here are the ones that come with practitioner case analysis. The evidence draws on a single practitioner's project record, without a controlled comparison against projects run on discipline-based tracking, and it stays within U.S. industrial and mission-critical contexts. A further limitation concerns scope: the case evidence in Section 5 covers turnover and commissioning readiness quality, testing, equipment, and turnover-documentation records the author directly verified and does not extend to verified operator-training, spare-parts, or maintenance-procedure records, even though the OR concept as defined in Section 1 spans those streams as well. The quantitative figures are used to convey scale and data structure, not to measure the framework's effect, which the present data cannot isolate. Broader, multi-practitioner data would be needed before generalizing to other delivery models and regulatory environments, and future work extending the case evidence to the training, spare-parts, and maintenance streams would strengthen the framework's claim to cover operational readiness in the fuller sense.

Where the work should head next is from document-based and spreadsheet-based turnover toward digital twin-integrated real-time operational readiness dashboards, with the system-based data structure described here serving as

the organizing schema for a facility's digital representation. The digital twin and semantic-enrichment research surveyed above offers a plausible representational foundation [6,7,10]; what remains is to bind that representation to the commissioning-system readiness logic startup actually requires, and to validate the resulting dashboards against startup outcomes on live projects. A readiness-oriented digital twin would do more than mirror a facility's physical state. It would answer, continuously and by system, the question mechanical completion leaves open: is this facility ready to be operated.

Several intermediate steps sit between current practice and that goal. Standardizing the system-based schema so a commissioning system, its streams, and their dependency logic are represented consistently across projects and contractors would be a prerequisite for any cross-project validation. Wiring the schema so sensing and progress-monitoring outputs feed it automatically, rather than by manual entry, would test whether the surveyed technologies can actually cut the labor of turnover instead of adding a parallel reporting burden. Comparing readiness-tracked projects against discipline-tracked baselines, where such baselines can be built, would start to quantify an effect this article can only describe. None of these steps demands a technology that does not already exist; each asks that the technology be organized around the per-system readiness verdict.

AUTHOR CONTRIBUTIONS

Rushi Jariwala is the sole author and is responsible for the conceptualization, methodology, formal analysis, and investigation underlying this manuscript, including the practitioner framework and the case analysis drawn from the author's project execution experience. Manuscript preparation, including drafting, editing, and formatting support, was assisted by a professional academic writing service engaged by the author; the author reviewed, verified, and approved all technical content, figures, and conclusions, and takes full responsibility for the accuracy and integrity of the manuscript.

CONFLICTS OF INTEREST

The author declares no conflicts of interest.

FUNDING

This research received no external funding.

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES

During the preparation of this manuscript, the author engaged a professional academic manuscript preparation service, which used generative AI tools to assist with drafting, language editing, and formatting based on the author's project records, notes, and directions. The author reviewed, verified, and edited all technical content, figures, and conclusions, and takes full responsibility for the content of this manuscript.

HIGHLIGHTS

- Operational readiness is treated as a planning-stage objective that governs the full project lifecycle, not a milestone inspected only when construction ends.
- Digital tracking and dashboard systems should organize completion and quality data by commissioning system, not by construction discipline, if they are to support system-level readiness verification.
- The framework is drawn from and illustrated against four U.S. capital projects spanning petrochemical, methanol, and modular data center construction.
- The work sets out data-structure and cadence requirements that digital twin, IoT, and progress-monitoring tools would need to meet to support near-real-time operational readiness visibility.

REFERENCES

1. Construction Industry Institute. Construction Readiness Assessment for Productivity Improvement. Final Report FR-DCC-02. Austin, TX: CII; 2018. <https://www.construction-institute.org/construction-readiness-assessment-for-productivity-improvement>
2. Ibrahim MW, Labib Y, Veeramani D, Hanna AS, Russell J. Comprehensive Model for Construction Readiness Assessment. *Journal of Management in Engineering*. 2021;37(1):04020088. [https://doi.org/10.1061/\(ASCE\)ME.1943-5479.0000832](https://doi.org/10.1061/(ASCE)ME.1943-5479.0000832)
3. Griego R, Leite F. Premature Construction Start Interruptions: How Awareness Could Prevent Disputes and Litigations. *Journal of Legal Affairs and Dispute Resolution in Engineering and Construction*. 2017;9(2):04516016. [https://doi.org/10.1061/\(ASCE\)LA.1943-4170.0000210](https://doi.org/10.1061/(ASCE)LA.1943-4170.0000210)
4. Uptime Institute. Annual Outage Analysis 2025. Uptime Intelligence, Keynote Report 173; 2025. <https://uptimeinstitute.com/resources/research-and-reports/annual-outage-analysis-2025>
5. Barbosa F, Woetzel J, Mischke J, Ribeirinho MJ, Sridhar M, Parsons M, Bertram N, Brown S. Reinventing Construction: A Route to Higher Productivity. McKinsey Global Institute; 2017. <https://www.mckinsey.com/~/media/mckinsey/business%20functions/operations/our%20insights/reinventing%20construction%20through%20a%20productivity%20revolution/mgi-reinventing-construction-a-route-to-higher-productivity-full-report.pdf>
6. Boje C, Guerriero A, Kubicki S, Rezgui Y. Towards a semantic Construction Digital Twin: directions for future research. *Automation in Construction*. 2020;114:103179. <https://doi.org/10.1016/j.autcon.2020.103179>
7. Deng M, Menassa CC, Kamat VR. From BIM to digital twins: a systematic review of the evolution of intelligent building representations in the AEC-FM industry. *Journal of Information Technology in Construction (ITcon)*. 2021;26:58-83. <https://doi.org/10.36680/j.itcon.2021.005>
8. Patacas J, Dawood N, Kassem M. BIM for facilities management: a framework and a common data environment using open standards. *Automation in Construction*. 2020;120:103366. <https://doi.org/10.1016/j.autcon.2020.103366>
9. Hosseini MR, Roelvink R, Papadonikolaki E, Edwards DJ, Pärn E. Integrating BIM into facility management: typology matrix of information handover requirements. *International Journal of Building Pathology and Adaptation*. 2018;36(1):2-14. <https://doi.org/10.1108/IJBPA-08-2017-0034>
10. Wang T, Gan VJL, Hu D, Liu H. Digital twin-enabled built environment sensing and monitoring through semantic enrichment of BIM with SensorML. *Automation in Construction*. 2022;144:104625. <https://doi.org/10.1016/j.autcon.2022.104625>
11. Zhao J, Pikas E, Seppänen O, Peltokorpi A. Using real-time indoor resource positioning to track the progress of tasks in construction sites. *Frontiers in Built Environment*. 2021;7:661166. <https://doi.org/10.3389/fbuil.2021.661166>
12. Wang J, Luo H, Peña-Mora F, Zhou W, Fang W. An integrated BIM-IoT framework for real-time quality monitoring in construction site. *Journal of Construction Engineering and Management*. 2024;150(11):05024013. <https://doi.org/10.1061/JCEMD4.COENG-14984>
13. Radman K, Jelodar MB, Lovreglio R, Ghazizadeh E, Wilkinson S. Real-time tracking and analysis in construction projects: a RealCONs framework. *Advanced Engineering Informatics*. 2025;67:103511. <https://doi.org/10.1016/j.aei.2025.103511>
14. Reja VK, Varghese K, Ha QP. Computer vision-based construction progress monitoring. *Automation in Construction*. 2022;138:104245. <https://doi.org/10.1016/j.autcon.2022.104245>
15. Lagos CI, Herrera RF, Mac Cawley AF, Alarcón LF. Predicting construction schedule performance with last planner system and machine learning. *Automation in Construction*. 2024;167:105716. <https://doi.org/10.1016/j.autcon.2024.105716>
16. Awada M, Srouf FJ, Srouf IM. Data-driven machine learning approach to integrate field submittals in project scheduling. *Journal of Management in Engineering*. 2021;37(1):04020104. [https://doi.org/10.1061/\(ASCE\)ME.1943-5479.0000873](https://doi.org/10.1061/(ASCE)ME.1943-5479.0000873)
17. Construction Industry Institute, Research Team 312. Critical Success Factors for Project Commissioning and Startup. Research Summary RS312-1. Austin, TX: CII; 2015. <https://www.construction-institute.org/critical-success-factors-for-project-commissioning-and-startup>
18. Alsakka F, Yu H, El-Chami I, Hamzeh F, Al-Hussein M. Digital twin for production estimation, scheduling and real-time monitoring in offsite construction. *Computers & Industrial Engineering*. 2024;191:110173. <https://doi.org/10.1016/j.cie.2024.110173>
19. Shin J, Choi B. Design and Implementation of Quality Information Management System for Modular Construction Factory. *Buildings*. 2022;12(5):654. <https://doi.org/10.3390/buildings12050654>
20. Sadeghi M, Elliott JW, Porro N, Strong K. Developing building information models (BIM) for building handover, operation and maintenance. *Journal of Facilities Management*. 2019;17(3):301-316. <https://doi.org/10.1108/JFM-04-2018-0029>
21. Abdelkarim SB, Ahmad AM, Naji K. A BIM-Based Framework for Managing Handover Information Loss. *Journal of Management in Engineering*. 2024;40(5):04024030. <https://doi.org/10.1061/JMENEA.MEENG-5623>
22. Xiao B, Xiao H, Wang J, Chen Y. Vision-based method for tracking workers by integrating deep learning instance segmentation in off-site construction. *Automation in Construction*. 2022;136:104148. <https://doi.org/10.1016/j.autcon.2022.104148>
23. Jariwala R. Integrated Turnover Management Systems for Mission-Critical Infrastructure Projects: Aligning Construction Completion with Operational Readiness. Manuscript submitted for publication; 2026.

24. Bhatnagar A, Bonagiri VRK. Operational Excellence through Cloud-Native Distributed Systems. *International Journal of Computational and Experimental Science and Engineering*. 2026;12(2). <https://doi.org/10.22399/ijcesen.5207>
25. Deshpande D, Yaramati KC, Bobba CT, Priya M, Pasupuleti SC, Gubbala SA, Thakkar N. Development, Validity, and Test Reliability and Reproducibility of Facial Palsy-Associated Facial Movement: A Cross-Sectional Study. *Journal of Pharmacy & Bioallied Sciences*. 2025;17(Suppl 4):S3120.
26. Gaggar H. Policy-Driven Access Control for Secure Deployment of Autonomous AI Agents in Enterprise Environments. *International Journal of Computational and Experimental Science and Engineering*. 2025;11(4). <https://doi.org/10.22399/ijcesen.5497>