

REFERENCE ELEVATION · NOT TO SCALE

THEORY, MEASUREMENT AND GOVERNANCE OF OPERATIONAL READINESS

The Operational Capability Gap

Jorge Luiz Dias

JLD Consultoria em Gestão de Projetos

Version 1.0 · 2026

ABSTRACT

Every capital project promises a capacity that operations must demonstrate. Between the capacity installed at the end of construction and the capacity actually proven in steady operation there is a distance — the operational capability gap. This paper defines that gap, proposes a way to measure it with evidence, and describes the governance required to close it before start-up turns technical decisions into emergency decisions.

Keywords: operational readiness; operational capability; capital projects; commissioning; evidence governance; risk management.

AUTHOR Jorge Luiz Dias
ORGANIZATION JLD Consultoria em Gestão de Projetos
VERSION 1.0, published in 2026
CONTACT contato@jorgeluizdias.com.br
RIGHTS All rights reserved. Reproduction permitted with attribution, without alteration of content.

CONTENTS

01 The problem nobody measures 3
02 What the operational capability gap is 3
03 The three laws 5
04 How it is measured 5
05 Governance: the rule that is not negotiable 7
06 The cost of intervening late 8
07 Application in seven steps 9
08 What the method is not 10
09 Evidence and limits 11
10 About the author 11
11 Frequently asked questions 12
References 13

01 The problem nobody measures

The project closes on time, the budget holds, construction ends. Months later, the operating result is still not what justified the investment.

The scene is familiar to anyone who follows capital projects in process industries. The asset is delivered, systems are energised, the first feed enters. Then begins a period in which the plant produces below plan, consumes more than plan and depends on external support to remain stable. That period is called ramp-up. It is treated as normal, as a foreseen stage, and often it is — provided one knows how long it should last and what must happen for it to end.

The problem is that, in most projects, nobody defined that criterion. The question “is operations ready?” is answered with a perception: the team is trained, procedures exist, commissioning has advanced. These are true and insufficient statements. They do not distinguish what was installed from what was demonstrated, nor what was documented from what was verified under real operating conditions.

The distinction the progress report does not make

A physical progress report measures what was built. A commissioning report measures what was tested. Neither measures what the organisation is able to operate in a sustained way, with the people who will actually be in the plant, the systems that will actually be available and the conditions that will actually exist at start-up. Installed capacity is a property of the asset. Operational capability is a property of the sociotechnical system that operates the asset.

That difference is not semantic. It explains why technically successful projects produce operations that take long to stabilise, and why the usual response — more training, more specialists, extended technical assistance — treats symptoms without touching the cause.

> THE QUESTION THIS DOCUMENT ANSWERS
>
> How to measure, with verifiable evidence, the distance between the capacity the project promised and the capacity operations demonstrates — and how to govern that distance before it turns into lost value.

The cost that does not appear in the business case

The business case compares investment with expected return. It rarely states how much of that return depends on a capability that does not yet exist and will be developed after start-up. When that dependency is not declared, the risk is not eliminated: it is transferred to operations, which must build capability while producing, without a budget for it.

The gap, therefore, is not only technical. It is a transfer of risk between two organisations with different objectives and deadlines — project and operations — often without either being aware of the value at stake.

The gap is not a metaphor. It has a unit, a measurement method and an economic consequence. That is what the following chapters define.

02 What the operational capability gap is

The gap is the difference between required capability and available capability demonstrated with evidence, at a given moment in the asset life cycle.

Three quantities must be distinguished precisely, because confusing them is the origin of the problem.

Required capability

What the business needs operations to deliver: nominal production, quality, availability, unit cost, safety and compliance requirements. It is defined in the business case and in project requirements. It is not an aspiration; it is a verifiable specification.

Installed capacity

What the asset can deliver when all its systems are available and operating within design parameters. It is a property of project and construction, proven by acceptance tests. A plant can have full installed capacity and low operational capability.

Demonstrated capability

What the organisation proves it can operate sustainably, with the people, procedures, systems and support that exist at that moment. It is not what is expected to happen at start-up; it is what has already been evidenced, verified and accepted by those accountable for operations.

> Gap = required capability – available demonstrated capability

The decisive word is **demonstrated**. A capability declared in a presentation, forecast in a schedule or estimated by analogy does not reduce the gap. Only capability that can be shown, on site, with traceable evidence, under the conditions in which operations will occur, reduces the gap.

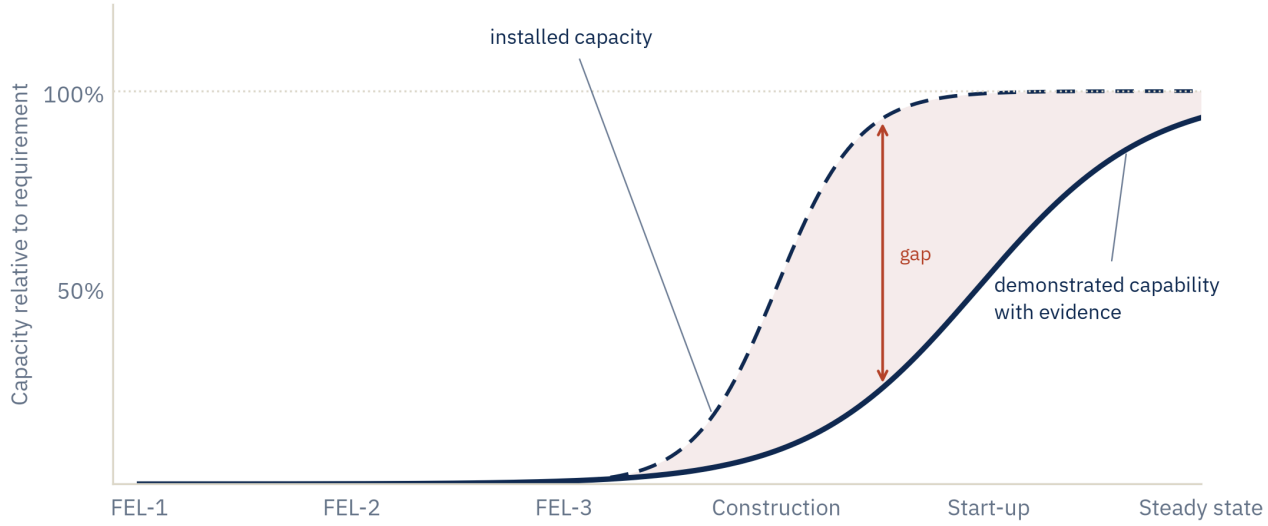


Figure 1: Installed capacity rises with the completion of construction; demonstrated capability rises with operation and evidence. The area between the curves is the gap. Illustrative curve, not to absolute scale.

Four properties that change management

The gap is dynamic. It closes with evidence and reopens with changes in scope, people, systems or boundary conditions. Measuring once, before start-up, produces a photograph that ages quickly.
The gap is multidimensional. A project may be ready in process and not in maintenance; ready in training and not in emergency management. The aggregate hides the blocker.
The gap is context-specific. The same plant, with the same workforce, can have different gaps depending on operating regime, supply market and turnover history.
The gap is economic. Every percentage point of capability not demonstrated at start-up has a cost: production not delivered, additional consumption, emergency contracting, extended technical assistance. That cost can be estimated before start-up — and that is what allows decisions on where to intervene.

What “demonstrated” excludes

It excludes capability that exists on paper and has not been tested; that was tested on a bench but not under operating conditions; that depends on a specialist who will be demobilised at start-up; and that depends on a system not yet integrated. Each of these exclusions is a measurement decision — and all of them push the index down, not up.

The gap across the life cycle

In definition, the gap is one of requirement: what operations will need and the project has not yet incorporated. In construction, it is one of dependency: what operations must receive and the project has not yet

delivered. In commissioning, it is one of evidence: what already works and has not yet been demonstrated. At start-up, it is one of performance: what was demonstrated and does not hold in steady operation. The same word describes four different problems, and each requires a different intervention.

03 The three laws

Three propositions support the method. They are not abstract principles: each has a direct consequence for what is measured and for when decisions are made.

Law I – Technical readiness is not operational readiness

A system can be technically complete and operations not ready to operate it. Technical completeness is measured against engineering specifications; operational readiness is measured against the ability to operate, maintain and recover the system with available resources. Confusing the two leads to declaring readiness at the moment when the evidence is still engineering evidence, not operating evidence.

Law II – The cost of closing the gap rises as start-up approaches

Closing a gap in definition costs a decision. Closing it during construction costs a change. Closing it in commissioning costs an intervention under schedule pressure. Closing it after start-up costs lost production, safety risk and credibility. The rise is not linear: it follows the growing number of decisions already taken that would have to be undone.

Law III – Operational capability is not commissioned at the end

Operational capability is built across the life cycle, with people, processes and systems prepared and verified step by step. Treating it as a closing activity is the error that produces long ramp-ups: at the end, there is no time left to develop competence, only to improvise.

> PRACTICAL CONSEQUENCE
>
> If the three laws hold, measuring the gap stops being a closing report and becomes an investment decision instrument. Measure early to decide early.

What the three laws imply together

The first law defines the object — operational capability, not technical completeness. The second defines time — the earlier, the lower the cost. The third defines the method — capability built across the cycle, not concentrated at the end. Together they move the question from “are we ready?” to “what is the gap, where is it and when will it be closed?”

04 How it is measured

Measuring the gap is not applying a questionnaire. It is establishing, for each domain, what evidence exists, what evidence is sufficient and which blocker prevents progression.

Measurement is organised in three layers: capability domains, evidence scale and weighted index. The three answer different questions and do not replace one another.

4.1 The six capability domains

The domains cover the sociotechnical system that operates the asset. They are the same in any sector; what changes is vocabulary and the relative weight of each.

DOMAIN	WHAT IS ASSESSED
D1	People and organisation: competence, staffing, roles, shift structure, operations leadership
D2	Operations and process: procedures, routines, parameters, operating performance, project interface
D3	Technology, data and systems: automation, integration, data reliability, support systems
D4	Maintenance, reliability and assets: plans, criticality, spares, integrity, availability
D5	Health, safety, legal requirements and emergencies: compliance, risk analysis, emergency response
D6	Documentation, supply and support: records, contracts, supply chain, operations support

Domains are not departments. A blocker in D4 may originate in a D1 decision; measurement must follow the operational effect, not the organisational chart.

4.2 The evidence scale

Each item is classified by the type of evidence that supports it. The scale is ordinal and deliberately conservative: it moves up a level only when verifiable evidence exists at the next level.

LEVEL	EVIDENCE	VERIFICATION
E0	None: the capability has not been addressed	Not applicable
E1	Declared: intent, plan or statement exists	Document without proven execution
E2	Documented: procedure, training or system exists	Record available, not tested in operation
E3	Verified: tested under representative conditions	Evidence independent of the executor
E4	Validated in operation: sustained performance in steady state	Operating series and operations sign-off

The scale does not measure effort. A training session delivered is E2; an operator who demonstrates the task under real conditions is E3; a team that sustains steady operation for a full cycle is E4.

4.3 The weighted index

Items are scored from 1 to 4 according to evidence, weighted by criticality and aggregated by domain and by project. The result is a readiness index, presented in bands that guide prioritisation — not release. The weights, sector calibration and band limits are proprietary and are not published in this document. What is published is the principle: the index exists to order interventions, and no index band authorises progression while a blocker remains open.

4.4 When to measure

Measurement is continuous, with formal milestones: in definition, to reveal scope gaps; in construction, to anticipate dependencies; in commissioning, to verify evidence; at start-up, to confirm capability; in stabilisation, to measure realised value. The minimum frequency is monthly during the critical phase.

> MEASUREMENT WITHOUT GOVERNANCE IS REPORTING
>
> An index that does not change decisions becomes bureaucracy within two cycles. The next chapter addresses what gives measurement consequence.

4.5 Common measurement errors

Measuring what is easy. Documents and training are easy to count; demonstrated capability is not. An index built on what is easy to measure produces an optimistic reading.

Accepting evidence from the executor. Verification loses value when the same team produces and validates the evidence.

Aggregating before examining. The consolidated number hides the critical domain; reading must start from the blocker, not the average.

Measuring once. The gap is dynamic; a single photograph serves diagnosis, not governance.

Confusing absence of evidence with absence of risk. What has not been measured is not zero; it is unknown.

4.6 From index to decision

The index does not decide. It orders. A decision occurs when a blocker is compared with the critical date, the cost of closure and the risk of not closing. That is where measurement connects to the next chapter: without a rule that gives consequence to the blocker, the index becomes a historical series with no effect.

05 Governance: the rule that is not negotiable

Governance of the gap comes down to one rule and three instruments that make it executable.

5.1 The blocker rule

No aggregate percentage, readiness index, schedule pressure, cost consideration or acceptance of residual risk authorises progression while a blocker remains open without a formally recorded decision. The rule applies to everyone, including project leadership.

> No percentage overrides a blocker.

The rule exists because the aggregate is where the gap hides. A project at 92% readiness may be at 100% in a peripheral domain and 40% in the domain that prevents start-up. The number reassures; the blocker decides.

5.2 Master Tracker

The Master Tracker is the single record of blockers, evidence, owners and dates. It is not an action list: it is the source of truth on what prevents progression. Each line has a named owner, a closure criterion and the evidence required. The tracker is reviewed on a fixed cadence and is the basis of any executive report.

5.3 Recorded decision

When a blocker cannot be closed before a critical date, the decision to proceed is recorded with four elements: who decides, with what information, which residual risk is accepted and what the containment plan is. The recorded decision does not eliminate the gap; it makes the risk explicit and attributable. It is the opposite of tacit exception.

5.4 Maturity

The organisation’s maturity to manage the gap is classified in five levels. The level is not a badge: it describes what the organisation is able to do without depending on individual heroism.

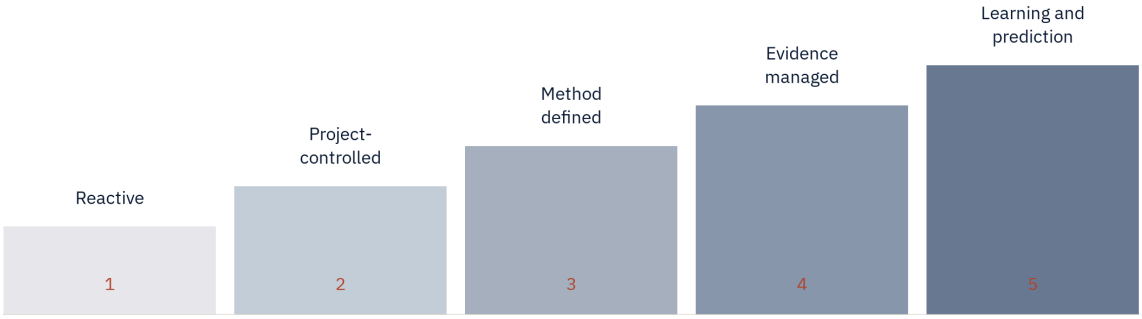


Figure 2: Five maturity levels. Progression requires evidence, not declaration. Qualitative scale, with no associated numeric score.

5.5 The cycle

Governance is a cycle, not a stage: the tracker feeds the decision, the decision requires evidence, the evidence updates the tracker. At each review, the gap is reopened and re-measured — including after start-up.

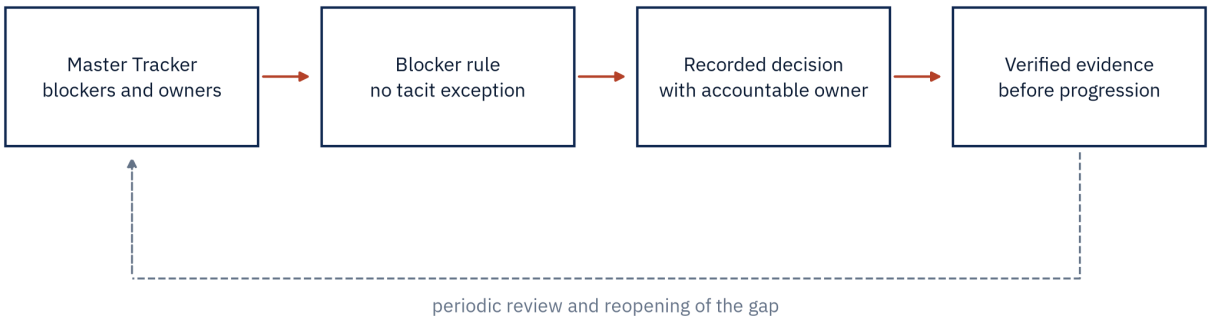


Figure 3: The governance cycle of the gap: record, rule, decision and evidence, with periodic review.

5.6 What happens without the rule

Without the rule, progression is decided by agenda availability and optimism. The blocker is recorded, then reclassified as “acceptable risk”, then transferred to operations. Each transfer looks reasonable in isolation; the set produces a technically authorised and operationally unprepared start-up.

The rule does not exist to prevent start-up. It exists to ensure that, if start-up occurs with an open blocker, the decision is explicit, attributable and documented — and that the cost of that choice is known to those with authority to assume it.

06 The cost of intervening late

The second law is not a moral warning. It is a cost relationship that can be estimated and used to decide when to intervene.

The cost of closing a gap rises with the number of decisions already taken that would have to be undone. In definition, closing a gap is adjusting a requirement. In construction, it is changing a contract. In commissioning, it is reorganising a sequence. At start-up, it is operating below plan while correcting what should already have been ready.

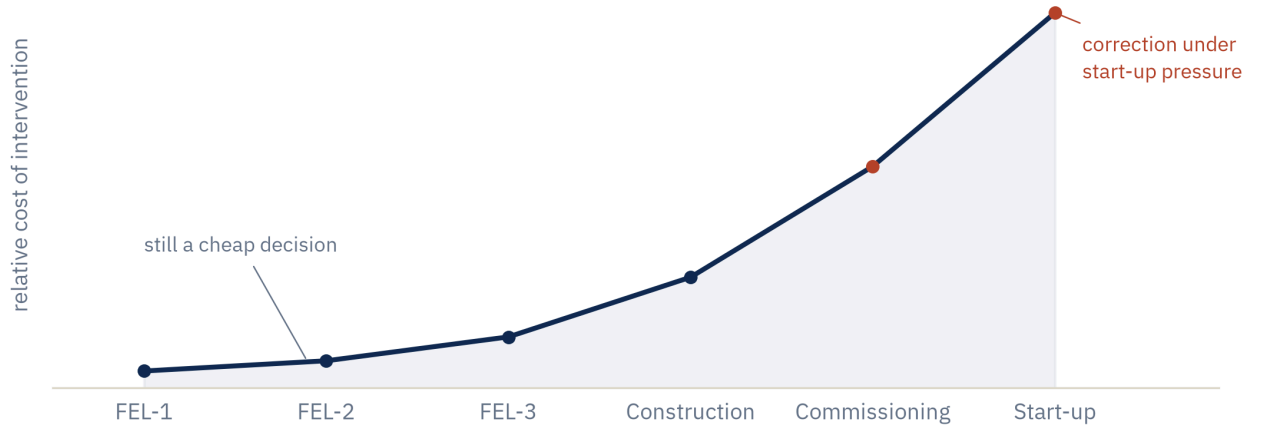


Figure 4: Relative cost of intervention by phase. The curve is illustrative and does not represent absolute values: what matters is its shape — cost grows faster than the time available to act.

Four factors explain the non-linearity:

- **Coupling.** The later the decision, the more decisions depend on it and the more expensive it is to reverse.
- **Attention capacity.** In commissioning the organisation is saturated; intervention competes with start-up.
- **Loss of options.** Before start-up there are alternatives. Afterwards there is only contingency.
- **Opportunity cost.** Every day of ramp-up has a value that is not recovered.

> THE DECISION WINDOW
>
> The objective is not to eliminate the gap before start-up at any cost. It is to decide, with information, where it pays to intervene early and where residual risk is acceptable. Measurement exists to make that choice explicit.

6.1 How to estimate

The estimate does not require accounting precision. It requires three inputs: the value of capability unavailable per day, the time needed to close the gap in each phase and the probability that the blocker materialises. With those three, it is possible to compare the cost of intervening now with the expected cost of intervening later — and to decide transparently.

6.2 The error of waiting for certainty

Waiting for certainty has a cost. When the decision is postponed until the blocker materialises, the organisation loses the cheapest alternative and operates only with contingency. The second law does not say every gap must be closed immediately; it says the cost of closing rises, and that the decision to postpone must be as explicit as the decision to intervene.

07 Application in seven steps

The seven steps are a management sequence, not project phases. They repeat in each decision cycle and integrate with investment gates.

	STEP	QUESTION IT ANSWERS	DELIVERABLE
01	Frame	What is the scope, the readiness criterion and who is accountable for it?	Readiness charter
02	Mobilise	Who does it, with what resources and under what governance?	Mobilisation plan
03	Plan	Which blockers exist and how will they be closed?	Master Tracker and closure plan
04	Execute	What is being done to close each gap?	Evidence per item
05	Verify	Does the evidence withstand independent verification?	Verification report
06	Transfer	Who assumes operations and with what acceptance?	Transfer record
07	Stabilise	Does capability hold in steady state?	Stabilisation curve and realised value

What each step delivers

Frame defines the readiness criterion before it is negotiated. **Mobilise** ensures the readiness function has authority to measure. **Plan** turns the gap into a blocker with an owner. **Execute** produces evidence, not activity. **Verify** separates what was done from what was proven. **Transfer** establishes who answers for performance. **Stabilise** measures the value the business actually captured.

Integration with the project cycle

In definition, steps 01 to 03 produce requirements and blockers before scope is frozen. In construction, 03 to 05 keep the tracker alive and verify evidence. In commissioning, 05 and 06 transfer operations with criteria. At start-up, 06 and 07 measure demonstrated capability and realised value.

Roles

Accountability for readiness belongs to operations, not the project. The project provides the asset and engineering evidence; operations accepts the transfer and answers for performance. The readiness function coordinates, measures and escalates — it replaces neither.

08 What the method is not

Clarity about what the method does not do is part of the method.

It is not a checklist

Checklists verify existence; the gap measures demonstrated capability. An item marked complete can remain at E1 if there is no evidence that the capability exists in operation.

It is not software

The Master Tracker can run on a spreadsheet. The tool does not create discipline; the blocker rule does.

It is not a guarantee of results

The method increases the probability that the gap is identified and addressed in time. It does not eliminate execution risk, nor does it replace engineering, commissioning or project management.

It does not replace technical commissioning

Commissioning verifies systems; operational readiness verifies the ability to operate, maintain and recover those systems. They are complementary disciplines with distinct technical authorities.

It is not a compliance audit

The objective is not to punish deviations but to reveal blockers while there is still time to decide. Measurement loses value when it becomes an instrument of blame.

09 Evidence and limits

A theory of operational readiness must be explicit about what has been observed and what remains a hypothesis.

Application

The method was formulated and applied in a large industrial expansion of approximately BRL 650 million and remains in use in an operating industrial complex, in change and capability management. Cases are presented anonymously under confidentiality commitments. No performance figures are presented because results depend on context, execution and factors external to the method.

Limits

Measurement quality depends on the quality of available evidence and on the readiness function's access to real decisions. In organisations where information is withheld, the index measures the information, not the capability. Sector calibration of weights requires history; without history, prioritisation is qualitative.

What still needs to be tested

The relationship between the gap measured at start-up and time to stabilisation is the central hypothesis to be verified with data from multiple projects. The theory predicts correlation; magnitude depends on sector, scale and maturity. That is the research agenda the method proposes.

How the theory would be falsified

If, across a relevant set of projects, the gap measured at start-up bore no relationship to time to stabilisation, the central hypothesis would lose strength. If applying the method changed no decisions and reduced no open blockers at start-up, the method would be only a report. The theory holds insofar as it produces different decisions — and it can be contested with data.

10 About the author



JORGE LUIZ DIAS

Jorge Luiz Dias

Creator of the Operational Capability Gap Theory

Jorge Luiz Dias has worked in industry since 1992, in capital projects and operations in the aluminium sector. He was production manager at an alumina refinery, in the Bayer process, and is today a major projects manager. His career combines operating plants in steady state with leading large-scale investments.

He holds a degree in Business Administration, a postgraduate qualification in Production Engineering and an MBA in Project Management. He formulated the Operational Capability Gap Theory from the observation that the distance between what a project delivers and what operations captures was not measured with the same rigour applied to schedule and cost.

CONTACT

contato@jorgeluizdias.com.br
jldias-ksheae3q.manus.space

11 Frequently asked questions

Is this not the same as operational readiness?

No. Operational readiness is a broad discipline rooted in commissioning, process safety and reliability. OCGT does not replace it: it defines a measurable quantity — the gap — and a governance rule that gives measurement consequence. Readiness frameworks describe what should exist; the theory proposes how to measure what actually exists and how to prevent progression from ignoring what is missing.

Does the method work for small projects?

Yes, proportionally. In a small project the apparatus reduces to a simple tracker, an evidence scale and a monthly review. The blocker rule does not change: it is what prevents schedule pressure from turning a pending item into a tacit exception.

Is software required?

No. The Master Tracker can run on a spreadsheet. The tool does not create discipline; the rule does. Software helps scale and audit, but does not replace the decision not to progress with an open blocker.

How to start in thirty days?

Three moves: define the readiness criterion and applicable domains; list known blockers and classify them by evidence; establish the review cadence and the body that decides. The rest is execution.

Who should lead?

Accountability for performance belongs to operations. The readiness function coordinates measurement, maintains the tracker and escalates blockers. The project provides engineering and commissioning evidence. When leadership is transferred to the project, measurement reflects what the project wants to show, not what operations needs to know.

How does this relate to commissioning?

They are complementary and not interchangeable. Commissioning verifies systems against technical specifications; operational readiness verifies the ability to operate, maintain and recover those systems with the people and resources available. A system can pass commissioning and fail in operation.

Are the index weights public?

No. The principle is public; sector calibration, weights and band limits are proprietary. Publishing calibration without the history that supports it would produce false precision — and the method would lose its main virtue, which is making judgement explicit.

References

Flyvbjerg, B. (2014). What You Should Know About Megaprojects and Why: An Overview. **Project Management Journal**, 45(2), 6–19.

Merrow, E. W. (2011). **Industrial Megaprojects: Concepts, Strategies, and Practices for Success**. Hoboken: Wiley.

Project Management Institute. (2021). **A Guide to the Project Management Body of Knowledge (PMBOK Guide)** (7th ed.). Newtown Square: PMI.

Reason, J. (1997). **Managing the Risks of Organizational Accidents**. Aldershot: Ashgate.

Vaughan, D. (1996). **The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA**. Chicago: University of Chicago Press.

Weick, K. E., & Sutcliffe, K. M. (2007). **Managing the Unexpected: Resilient Performance in an Age of Uncertainty**. San Francisco: Jossey-Bass.

International Atomic Energy Agency. (2014). **Commissioning for Nuclear Power Plants**. IAEA Safety Standards Series. Vienna: IAEA.

The references indicate the conceptual basis consulted. The Operational Capability Gap Theory is an original formulation by the author and is not presented by these works.

APPLICATION

The gap exists in every project. The difference is whether it will be measured.

INITIAL DIAGNOSIS

A sixty-minute conversation applied to a specific project or portfolio, to identify the current gap, the known blockers and what must be decided before start-up.

CONTACT

contato@jorgeluizdias.com.br
jldias-ksheae3q.manus.space

JORGE LUIZ DIAS · JLD CONSULTORIA EM GESTÃO DE PROJETOS