ARTIFICIAL INTELLIGENCE: GOVERNING THROUGH CLARITY

The Navigation Manual for Those Who Refuse to Get Lost in the Fog

Yvan Cognasse & Paul Neuman

ITSD

GOVERNING THROUGH CLARITY

About This Extract

AI: Governing Through Clarity is a practical governance manual for organizations trying to move from AI experimentation to accountable execution.

Its central idea is simple: most organizations do not lack AI ambition, tools, or pilots. What they often lack is clarity around decision rights, ownership, accountability, data readiness, risk, and the discipline to decide what should be scaled, challenged, or stopped.

This manual is not a technical catalogue. It is written for executives, managers, transformation leaders, AI sponsors, consultants, and teams who need to make AI governable, responsible, and executable.

The full book is available in English, French, and German, in both paperback and eBook formats.

- English: https://relinks.me/B0GZ8DR6FM
- French: https://relinks.me/B0GZ7YKKHN
- German: https://relinks.me/B0GZ8B6HC2

For reviews, executive discussions, roundtables, practitioner sessions, or interviews:

Paul Neuman
e-mail: paul.neuman@itsd.pro / LinkedIn: paulneuman

Yvan Cognasse
e-mail: yvan.cognasse@gmail.com / LinkedIn: yvancognasse

ARTIFICIAL INTELLIGENCE: GOVERNING THROUGH CLARITY

The navigation manual for those who refuse to get lost in the fog

The information contained in this manual is provided for general guidance to support the governance, deployment, and use of artificial intelligence within organizations.

First edition, May 2026

Published by ITSD / IT Services Demystified Editions

ISBN: 979-8-259-12882-8

Contents

Foreword

We each have more than thirty years of field experience behind us. Thirty years spent building, deploying, and sometimes repairing what technology promises organizations.

We have implemented ERP systems in dozens of countries, and we have seen what they brought at their best: robustness, data harmonization, and a backbone for companies that had become global. We supported the rise of Big Data, which opened up immense possibilities for customer insight. The Cloud architectures we built gave organizations a degree of agility they could not have imagined ten years earlier. Each wave laid down another layer of foundation. None delivered on every promise, but each left something solid behind.

AI is arriving on top of those foundations. That is what makes the current moment so decisive. For the first time, technology can draw on everything that was built before it, provided the organization is ready to receive it.

Yet what we observe, project after project, is a widening gap between the power of the tool and the maturity of those deploying it. Considerable budgets disappear into prototypes with no future. Brilliant teams burn out under contradictory injunctions. Sincere executives find themselves caught between market pressure and the inability to measure anything at all.

So the organization starts moving frantically. It launches prototypes, hires consulting firms, and opens AI labs. It creates the appearance of movement without producing any real movement at all.

Behind this theater of innovation, reality is harsher. Projects stall, then quietly die out. Data resists. Responsibilities withdraw into the fog. This is not a spectacular failure, but a slow, costly, and silent drift.

That is what we call the Fog.

The fog is not the failure of a technology, because AI works. It is the massive waste of resources, time, and credibility. It is the confusion of speed with direction, demonstration with transformation, and content production with value creation.

In an aligned organization, AI amplifies clarity. In a confused organization, it amplifies chaos. What AI ultimately exposes is not a technological weakness, but a governance weakness: an inability to decide, a reluctance to name an owner, and a long-standing tolerance for ambiguity that has become untenable.

That fog now meets a particularly deceptive technology: AI does not merely reveal flaws; it dresses them up. It clothes approximation in impeccable syntax and gives disorder the appearance of mastery.

The illusion is dangerous:

- Because the output is fluid, we believe it is true;
- Because the sentence is well written, we lend it authority;
- Because the tone is confident, we assume it is under control.

By shifting trust from verification to appearance, the organization takes on total risk: managerial, legal, and reputational.

This manual rests on one conviction: AI does not need a story; it needs a framework. It requires trade-offs, healthy data, and assumed responsibilities. Without those, it merely amplifies the fragilities the organization has not yet found the time, or the courage, to address.

Today, courage does not mean moving faster. It means becoming clearer.

- Stopping what does not work, even when a powerful sponsor stands behind the project;

- Setting minimum conditions before investing;
- Treating data as an asset, not as waste.

What we are arguing for here is innovation that is not merely fast, but useful, controlled, and measurable. William Gibson reminded us that the future is already here, just unevenly distributed; this manual is meant to help those who refuse to suffer it blindly navigate it with greater lucidity.

We hope you find this manual useful. And above all, hold your course.

We only cure a weakness
once we recognize it.

François de La Rochefoucauld, Maxims (1665)

Introduction

This manual is written for executives, executive committee members, and leaders confronted with the sudden arrival of AI. It is intended for those who want to act with lucidity, responsibility, and effectiveness.

If you are looking for a technical demonstration, a catalog of tools, or promises of miraculous gains, this manual is not for you. It is for those who must decide, commit budgets, assume risks, and produce measurable outcomes.

It is structured in four parts: The Lucidity Shock, The Truth Diagnosis, The Clarity Method, and The Embodied Leadership. Each part stands on its own and can be read independently according to your immediate priorities.

You can approach this work along two reading paths:

1. Linear Path: a complete journey, from recognizing the drift to applying a structured method for governing and deploying AI. Ideal if your projects are stalling or your teams are burning out;
2. Diagnostic Path: an issue-based approach. Each part can be consulted independently to address a specific urgent challenge.

By the end of this manual, you will not merely know what AI is. More importantly, you will know how to:

- Identify the minimum conditions before investing another euro;
- Distinguish a value-creating project from a reckless one;
- Build governance capable of making clear trade-offs;
- Treat data as a genuine strategic asset;
- Decide whether to scale or stop, with full awareness of the consequences.

This manual does not replace your technical experts or your legal counsel. It provides a decision framework.

AI is no longer limited to helping produce text, analyzis, or synthesis. In more and more situations, it is beginning to chain tasks together, query systems, and trigger actions. The more it acts, the more governance becomes a central issue.

Its purpose is to help you ask the right questions at the right time, so that market pressure or technological appearance does not replace your judgment.

Finally, this is neither an indictment of artificial intelligence nor a technophile plea on its behalf. It is based on a simple observation: AI is not the problem in itself; the problem lies in how it is introduced, governed, and deployed within organizations.

If it begins with failures, drifts, and absurd situations, it is not to discourage action, but to prevent blind action. The shock is not a conclusion. It is a necessary step that makes diagnosis, method, and ultimately disciplined execution possible.

AI is a powerful amplifier. The challenge is to ensure it amplifies in the right direction.

Strip them of the legend they wrap themselves in; look at them as they are.

Marcus Aurelius, Meditations (c. 170-180 AD)

PART 1
The Lucidity Shock

This first part is not here to dramatize, nor to lecture. It is here to look at the situation as it is. AI is arriving in a world that was already unstable, saturated with promises, slogans, and deferred decisions. It is moving too fast to leave us time to hide the cracks beneath another layer of communication.

We will not talk about solutions yet. We will map the terrain instead: the visible and invisible symptoms that derail projects the moment they collide with reality. The objective is simple: to establish a lucid and unsparing reading framework. Because before hoping to succeed with AI, you first need to know what ground you are standing on.

Every technological revolution begins with a honeymoon phase. AI is no exception; if anything, it is going through one at accelerated speed.

The promises are not merely flowing in; they are saturating the space. We are sold the end of repetitive work, infinite creativity, and the instant resolution of long-standing problems. The demonstrations are dazzling: in three clicks, a video generates itself, code writes itself, and a report summarizes itself.

Faced with that apparent magic, the temptation is immense to suspend critical judgment. We want to believe that signing a cheque and adopting the right tools will make the complexity of reality evaporate and finally bring everything into alignment.

And yet, if you have more than twenty years of professional memory, a strong sense of déjà vu becomes impossible to ignore:

- The ERP Era: it delivered, bringing robustness and often indispensable data harmonization, but sometimes at the price of considerable costs, permanent compromises, and a rigidity organizations still pay for today;
- The Big Data Era: we were promised an automatic goldmine. Too often, we ended up with swamps of unusable data;
- The Cloud Era: the promise was absolute agility. But we discovered that migrating servers was not enough to migrate culture.

Today, AI is replaying the same pattern: multiple initiatives, serial POCs, appealing prototypes, and then the wall of reality when integration, security, and maintenance become unavoidable.

The difference this time is speed. AI does not forgive vagueness. It does not compensate for organizational weaknesses; it exacerbates them. It acts like a magnifying mirror:

- If your processes are slow, it slows them down further by adding complexity;
- If your data is contradictory, it multiplies those inconsistencies at industrial scale;
- If your teams improvise, it pushes them into increasingly hazardous decisions.

What if the real danger were not technological? In most of the cases we observe, the decisive factor is managerial clarity; or its absence.

The Anatomy of the Fog

AI does not create chaos. It tends instead to expose what had long been waiting to be addressed.

It enters organizations where compromises have accumulated and where people have become accustomed to operating in approximation. The rigidity of tools is compensated for by human flexibility. Data errors are corrected manually. Overly heavy processes are bypassed with phone calls.

AI cannot cheat the system. It cannot make arrangements with reality. It requires explicit rules, clean data, and formalized processes.

When you plug AI into a vague organization, it does not make it smoother. It stalls, or worse, it automates the blur at high speed. That collision between the rigidity of the machine and the improvisation of the organization is what generates this confusion.

What Everyone Can See

Some signs of this opacity are staring everyone in the face in meetings, yet we rarely take the time to read them as weak signals.

What strikes first is the glaring contrast between the excitement of launch and the near-total inability to measure any tangible result a few months later. The organization seems seized by a frenzy of activity that barely conceals structural immobility.

It is the reign of Busy-ness, activity for activity's sake, over Business, measurable value. Calendars fill up with AI committees; inboxes overflow with internal newsletters about transformation. Yet the real performance indicators, margin, customer satisfaction, and cycle time, do not move by a millimetre. Behind that stagnation, the same drifts always return:

- Expeditions Without a Compass: initiatives are launched, tested, and announced. Sprints, hackathons, and strategic partnerships multiply. But no one really knows what must succeed, or how to measure it. Activity replaces strategy. Agitation is mistaken for progress;

- Institutional Procrastination: arbitration decides nothing. Decisions are postponed in the name of prudence, another study, a competitive benchmark, or one more committee. The organization spins in circles, trapped in an endless loop of validation while waiting for a consensus that will never come;
- Organizational Amnesia: with AI, organizations repeat the same mistakes they made with the Internet of Things or Blockchain, convinced that this time is different because the technology is brighter. Technology changes, but the human obstacles remain the same. Past failures are not capitalized on; they are erased so they can be repeated.

What the Slides Hide

The real chaos lies in the deeper layers, where PowerPoint presentations never linger. It is the hidden side of the technological iceberg.

While the executive committee gets excited by theoretical adoption curves, the field experiences another reality, rougher and far more corrosive.

- Technical Debt Explodes Silently: every quick POC adds another layer of undocumented complexity. Temporary bridges become permanent, weakening the entire information system for years to come. The future is mortgaged for short-term wins;
- Expert Fatigue Sets In: the best people are pulled onto ten priority projects at once. They are asked to build the future without being relieved of the present. The result is predictable: they burn out, disengage, or leave, taking critical knowledge with them;

- Cynicism Takes Root: after watching revolutions announced every six months that never see the next day, operational teams develop immunity to change. They no longer listen. They smile politely and wait for it to pass, just as they did with Blockchain or the Metaverse.

The Ghost Success of Predictive AI

An industrial leader celebrates in the executive committee a maintenance AI presented as 98% reliable. On the slides, it is a triumph. Down on the shop floor, the computer has been unplugged and is gathering dust. The shift manager explains: "*It kept sending false alerts every Tuesday because it confused machine cleaning with a breakdown.*"

The shock: the executive committee's success was a collective hallucination. No one had cleaned the data or explained the business context to the algorithm.

The Multiple Truths

Companies often think they lack computing power or sophisticated algorithms. But the question deserves to be turned around: what if the first thing they lack is a shared truth about the data itself?

Survival Glossary

AI-LIBI /ˌeɪ aɪ ˈlɪbaɪ/ n. (rhet. and iron.) the convenient use of the algorithm as a political shield, as in *"AI decided"*, in order to justify unpopular choices and relieve leadership of the burden of accountability.

AI GURU /ˌeɪ aɪ ˈgʊruː/ n. (iron.) a technological prophet selling absolute certainty about the replacement of humans, usually to conceal a striking lack of real operational experience.

AUTOPILOT /ˈɔːtəʊˌpaɪlət/ n. the dangerous belief in fully autonomous AI. Reminder: AI is a co-pilot; without human supervision, it drifts and eventually ends up in the ditch.

BUSY-NESS /ˈbɪzi nəs/ n. (iron. anglicism.) the art of looking busy, through committees, slides, and POCs, in order to simulate movement without producing economic value. Ant. *Business* (measurable results). The risk is that AI amplifies the former at the expense of the latter.

CLARITY PACT /ˈklærəti pækt/ n. the agreement required to stop organizational denial and address structural problems before launching technological innovation. Not a decorative charter, but a last attempt to reintroduce seriousness into the conversation.

COLLECTIVE HALLUCINATION /kəˌlektɪv həˌluːsɪˈneɪʃən/ n. a group phenomenon in which enthusiasm for a demo or POC eclipses technical warning signs and total unreadiness for production.

CONFIRMATION BIAS /ˌkɒnfəˈmeɪʃən ˈbaɪəs/ n. AI's natural tendency to validate a user's assumptions rather than challenge them, thereby supplying what looks like scientific endorsement for staying comfortably trapped in error.

About the Authors

Yvan Cognasse, a telecommunications engineer and graduate of École Centrale Paris, does not merely talk about AI. He turns it into measurable performance.

With thirty years of experience at the heart of major technology shifts, he now leads a team of AI specialists within a global technology leader. His role is to align strategic vision, technical architecture, and real business impact. He operates where it matters most: growth, financial performance, and governance.

As a director and facilitator of continuing education programs in AI, he passes on this same demand for clarity to future leaders. His signature is simple: he refuses ambiguity and builds AI frameworks that work, deliver, and endure.

Paul Neuman, a computer engineer and EPITA graduate with additional executive education at INSEAD, does not talk about results. He delivers them.

With thirty years on the front line of critical transformations and complex contractual issues, he now leads IT Services Demystified for international groups. His field of action is to untangle complexity, clarify strategic decisions, and secure operational ROI. As a trainer in digital governance for executives, he brings an uncompromising standard: AI must serve the business, never parasitize it.

He designs economic models, commercial offerings, and governance frameworks that enable teams to act faster and more accurately. His hallmark is radical pragmatism that serves frameworks that work, create measurable value, and last.

Budgets swallowed by dead-end prototypes. Brilliant teams exhausted by contradictory demands. Sincere leaders unable to measure what truly creates value. And everywhere, the same confusion: because AI impresses, we think we have mastered it.

AI masters nothing. It amplifies.

In a confused organization, it industrializes chaos. In a clear organization, it accelerates performance.

This manual promises neither revolution nor miracle. It delivers what few AI discussions offer today: **a framework** to distinguish value-creating projects from financial black holes disguised as innovation, identify the invisible debts that derail initiatives before they even reach production, and turn lucid abandonment into competitive advantage.

Written by two field practitioners, this book speaks to executives, business leaders, and technology decision-makers who refuse to navigate blindly.

AI is a magnifying mirror.
What it reveals must still be worth amplifying.

Yvan Cognasse is a telecommunications engineer and graduate of École Centrale de Paris. Leading a team of AI specialists at a global technology leader, he harnesses artificial intelligence for measurable performance and guides executives through these transformations. His guiding principles are threefold: cut through the fog, build robust frameworks, and make AI a lever that works, delivers, and lasts.

Paul Neuman is a computer engineer, EPITA graduate, and INSEAD-trained executive. He supports international groups through complex and sensitive transformations. His expertise lies in making complex situations clearer, reducing uncertainty, securing operational results, and aligning technology with business priorities. Through IT Services Demystified, he champions a simple credo: clarify AI so it serves real value.