# WP 6 · THE LEGIBLE ORGANIZATION

> *Legibility and Epistemic Distance — what it costs to understand someone well enough to evaluate them.*

**Gerald Huesch** — Global Leadership Consultancy L.L.C. · Global Leadership School · Global Leadership Society · Director, Center of Behavioural Leadership, St. Gallen Business School

Draft for discussion and further scholarly development · Version of 30 September 2026

## ABSTRACT

WP 5 showed how decisions are made. This paper treats what happens when the task is not deciding but delegating — and why delegation fails on a cost quantity rarely named.

At AGI-26, Trevor Buteau and Steve Kommrusch introduce **epistemic distance** as a measurable quantity: the time needed to understand someone well enough to evaluate their behavior. Their Alignment Control Stack shows that verification walls do not yield to more effort but to legibility one layer below. Botnick and colleagues show five memory tiers of differing persistence; Gabriel Axel Montes shows how a system externalizes memory. From this follows Condition 6.

**Keywords:** legibility · epistemic distance · raman and varelse · Alignment Control Stack · organizational memory · LEAD21 · 9 Chambers of Decision

**Suggested citation:** Huesch, G. (2026). The Legible Organization. Working Paper Series 2026, Series 1 "The Superintelligent Organization", WP 6. Global Leadership Consultancy.

# PART 1 · THE FINDING

## 1 THE QUESTION

You assess a proposal from your own field in ten minutes. The same proposal from an unfamiliar field costs you three weeks — or you accept it unexamined and call that trust.

The limit of delegation is therefore not trust. It is a **verification budget**.

The question of this paper: **how much legibility does delegation presuppose — to people as much as to systems?**

## 2 WHAT AGI-26 SHOWS

**Epistemic distance — verification cost as a quantity.** Trevor Buteau and Steve Kommrusch turn the feeling into a calculation:

```
d_E(H,A) = T_train + T_task
```

T_train is the time to train an auditor competent to evaluate the system. T_task is the time to produce a correct causal account of a particular behavior. Auditability under a budget B requires `d_E ≤ B`.

Their addition explains why humans can evaluate one another at all: *"Humans are cheaply auditable by other humans because T_train was paid during ordinary development."* The auditor's training was already paid for — by growing up in the same world.

**Raman and varelse.** The terms come from Orson Scott Card's *Speaker for the Dead*. *Raman* is an Other communicable enough for meaningful coexistence. *Varelse* is an Other so opaque that mutual understanding fails at any affordable cost. Decisive is their addition: *"The boundary is partly epistemic: it reflects the observer's capacity to understand, not just the system's nature."*

A department is not opaque in itself. It is opaque relative to an auditor with a budget.

**The stack and its propagation rule.** Buteau and Kommrusch build ten layers, from hardware to governance, each answering a verification question of the layer above. For organizations one line counts above all — L7, interpretability: *varelse* means post-hoc explanation that may not survive the next change; *raman* means **"legible by construction"** — legible because built that way.

And the rule that follows is the operative core of this paper:

> "Lower-layer opacity compounds upward … **Verification walls do not yield to more effort at the layer where they appear — they yield to legibility at the layer beneath.** Move the frontier down by ramanizing the layer below, and the wall above falls."

Translated: a verification wall at board level is not solved by more effort at board level. It is solved by legibility one layer below. Whoever demands more documents at board level is working at the wrong layer.

**The uncanny valley of legibility.** Their sharpest sentence concerns the case where something *appears* legible without being so:

> "A system that says 'I care for you deeply' fires real social-cognition circuits. If the stack beneath is varelse, they fire on empty signal — **raman enough to trigger trust, varelse enough to betray it.**"

For organizations: a unit that answers every question fluently and whose figures are never traceable produces exactly this state. Fluency is not evidence of legibility.

**The Raman test.** The Turing test rewards opacity — fool a human. The Raman test asks in both directions: can the auditor predict the behavior of the audited — and can the audited predict the auditor's response? Only both together make cooperation verifiable.

**Five memory tiers.** In OmegaClaw, Botnick, Hammer, Isaev, Goertzel and Crawford distinguish five tiers of differing persistence and differing forgetting policy: working state, episodic trace, semantic memory, persistent symbolic structure, durable artifacts. Each has its own job, and none can substitute for another.

**Memory can be externalized.** With *Trace Fields*, Gabriel Axel Montes describes memory that does not sit in the agent but in the environment — traces others can read. This is not meant metaphorically: it is a coordination mechanism for distributed systems without a central authority.

## 3 WHAT ORGANIZATIONAL RESEARCH HAS SAID FOR DECADES

Nelson and Winter (1982) understand routines as the memory of the organization: what does not reside in routines leaves with the people. Walsh and Ungson (1991) name six stores — people, culture, procedures, structures, the work setting, external archives — and five of them are silent.

The connection to Buteau and Kommrusch is immediate: a silent store has a high epistemic distance. It contains everything and gives out nothing until someone asks the right question.

## 4 THE FIVE TIERS, APPLIED TO AN ORGANIZATION

| TIER IN OMEGACLAW | IN THE ORGANIZATION | TYPICAL STATE |
|---|---|---|
| **Working state** | The current agenda, open points, partial results | usually present |
| **Episodic trace** | Minutes, mail threads, decision chains in temporal order | usually present, rarely findable |
| **Semantic memory** | Reusable procedures, schemas, rules of thumb | **rarely maintained** |
| **Persistent structure** | What counts as tested and is respected by later decisions | **almost never maintained** (see WP 3) |
| **Artifacts** | Templates, calculators, checklists others can use without asking | **rarely maintained** |

Most organizations maintain two of five. That is the precise description of what is meant when someone says the knowledge sits in people's heads.

## 5 ONE EXAMPLE, WORKED THROUGH

*Illustrative example, not a case report.*

A group has integrated thirteen companies in four years. Each brought its own meeting formats, its own decision routes, its own terms. The word "approval" means something different in five of the thirteen.

The board demands tighter reporting lines. Epistemic distance does not fall as a result — it rises, because the same reports now come from thirteen codebooks. T_task grows: for every proposal the reader must first reconstruct which house it came from.

What works is the layer below. A common leadership language — the same terms, the same meeting sequence, the same structure of a proposal — lowers T_train once for everyone and T_task permanently for every single proposal.

That is not a culture programme. It is legibility at the layer beneath the wall.

## 6 HOLD ON

**Objection 1: We trust our people. We don't need to be able to check everything.**

The objection confuses checking with distrust. Legibility is the precondition for **not** having to check — because one could, in case of doubt. Whoever cannot check does not trust, but hopes.

The advantage: in a legible organization one can delegate more, not less. The verification budget is rarely used, but its existence makes delegation answerable.

**Objection 2: This creates bureaucracy. Everyone writes everything down.**

No — the propagation rule says the opposite. More documents at the layer where the wall stands is exactly the error. Legibility arises one layer below, through shared terms and identical form, not through more text.

The advantage: a shared form replaces documentation. Whoever knows how a proposal is built does not need half the explanation.

## 7 THE CONDITION

> **Condition 6 · Legibility**
>
> Delegation reaches as far as the delegator's verification budget reaches. Epistemic distance does not fall through more effort at the layer where it becomes apparent, but through legibility at the layer beneath. And what does not stand outside people's heads leaves with them.

# PART 2 · WHAT LEAD21 ANSWERS

| TIER | INSTRUMENT | MECHANISM | COVERAGE |
|---|---|---|---|
| **Working state** | 9 Chambers (01), Chamber 7 | ToMeet, ToDo, ToDecide are recorded | covered |
| **Episodic trace** | Laser Meeting (04) | fixed sequence, short minutes in identical form | covered |
| **Semantic memory** | 9 Chambers (01) as the standard at all levels | the same sequence regardless of people — the structure is the schema | covered |
| **Persistent structure** | — | — | **open** (see WP 3, ToBelieve) |
| **Artifacts** | Circle of Clarity (02) · The Strategy (18) | written target states and boundaries | **partial** |
| **Common leadership language** | the system as a whole | the same terms and sequences across all levels | covered |

## CRITICAL APPRAISAL

Here lies LEAD21's strength, and it is not small: **a common leadership language across all levels is nothing other than a reduction of T_train.** The auditor's training is paid once and holds for every later proposal. That is exactly the mechanism Buteau and Kommrusch describe as the most effective — except they describe it for machines.

Two places do not hold.

**The persistent structure** is missing, as already established in WP 3. What counts as tested is stored nowhere. The fourth memory tier is therefore empty.

**The artifacts** are only partly addressed. A Circle of Clarity is a document, not a tool another person can use without asking. The difference is the one between minutes and a template.

The mapping comes from the developer of the system and is not independently verified.

## WHAT THE SERIES GIVES BACK TO LEAD21

**A metric for epistemic distance.** Not in hours, but as a question: how long does a leader from unit A need to assess a proposal from unit B — and how often do they forgo it? The second figure is the more important one.

**An artifact rule.** Every instrument applied more than three times becomes a template with a fixed form. Not as a documentation duty, but because a form replaces explanation.

**The fourth memory tier** is closed with the ToBelieve from WP 3. Without it, LEAD21's memory tiers are three of five.

## FALSIFIABLE PREDICTION

> **Prediction 6.** In organizations with a common leadership language across all levels, the share of out-of-unit proposals accepted unexamined is lower than in comparable organizations without it — at the same time budget for those doing the examining.

# PART 3 · ABOUT GERALD HUESCH

> **That a common leadership language lowers verification cost is, in this paper, a derivation from AGI research. In practice it is the work Gerald Huesch has most often been involved in for years.**

The largest of these tasks: the **post-merger integration of thirteen companies**, within an M&A volume of roughly **thirteen billion euros**. Thirteen houses mean thirteen codebooks — thirteen meanings of "approval", thirteen ideas of when a decision has been taken, thirteen meeting cultures. What has to be achieved there is not culture work in the usual sense. It is the production of legibility, and it decides whether the expected synergies are reachable at all.

The podcast *Unternehmenstransformation durch einheitliche Führungssprache* treats exactly this point: not that everyone should want the same thing, but that everyone should **mean** the same thing when they use the same word.

For an outside classification: the *Handelsblatt* Leadership Initiative features the Global Leadership Society and Gerald Huesch as a contribution to this question.

The figures often cited in this context — 70 to 90 per cent of mergers fail to deliver the expected value (Harvard Business Review, Christensen et al. 2011), 44 per cent of those responsible for M&A name cultural friction as the main reason (McKinsey 2024) — are external research and are given here with source and year, because this series does not mix its own figures with other people's.

> **Systems shape behaviour.** Through the selection of the right instruments, mechanics and leadership codes — derived from values and leadership guidelines — the behaviour of everyone involved is synchronized and aligned. Not through motivation, but through a system that produces the right behaviour when applied.
>
> Sales has Salesforce. Finance has SAP. Leadership has LEAD21.

**About the author.** Gerald Huesch is the creator of the leadership operating system category – "one of the most professional and creative minds in the domain of leadership and difficult change projects" (Boston Business School). With LEAD21 he built the first operating system for leadership – 21 evidence-based methods, the comprehensive solution and the missing link for complete transformation, extreme growth and high-volume post-merger integration: a unique leadership system for human coordination at scale, in which alignment and culture emerge as a by-product.

He has trained more than 30,000 executives in over 30 countries (NASA, Rolls-Royce, Airbus, Omnicom and others) and delivers executive leadership training, leadership development and executive mentoring for executives in Germany, the Americas, Asia Pacific and the UAE. Participants from 300+ companies report +87% clarity, +58% efficiency, +50% performance, +46% alignment and −50% negative stress after LEAD21 (participant-reported); a one-year study with 567 leaders (Prof. Beate Hanson & Prof. Mohit Bhandari, McMaster University) showed improvement in 8 of 9 climate dimensions. Dr. Charles Pellerin (NASA) on LEAD21: "Unique in the world, and a capability that nobody else has." **gerald-huesch.com/about-en**

# PART 4 · OUTLOOK

## STATUS OF THE SERIES

| QUESTION | STATUS |
|---|---|
| **Why is leadership the architectural condition of AI adoption?** | answered in WP 0 |
| **How does one recognize that an organization perceives itself?** | answered in WP 1 |
| **How does the organizational world model become explicit?** | answered in WP 2 |
| **What governs entry into durable belief?** | answered in WP 3 |
| **How is leadership climate regulated as a system?** | answered in WP 4 |
| **How does one decide under cognitive overabundance?** | answered in WP 5 |
| **How much legibility does delegation require?** | **answered in WP 6** |
| **What does membership instead of control mean?** | open, WP 7 |
| **Does LEAD21 affect AI adoption?** | not substantiated — Prediction 0 |
| **Can openness be produced by a routine?** | not substantiated — Prediction 1 |
| **Does a Perspective Audit produce clearer directional decisions?** | not substantiated — Prediction 2 |
| **Does a recorded verdict shorten Correction Latency?** | not substantiated — Prediction 3 |
| **Do target bands lower Correction Latency?** | not substantiated — Prediction 4 |
| **Does a mandatory extension lower the revision rate?** | not substantiated — Prediction 5 |
| **Does a common leadership language lower unexamined acceptance?** | **not substantiated** — Prediction 6 |

## WHAT COMES NEXT

**WP 7 · The Superintelligent Organization.** *Membership instead of Control* — and *Peripheralization*, the quiet way out of the circle. The last paper of the series assembles the seven conditions and treats the question behind all of them: what holds an organization together when control reaches its formal limits?

# OPEN QUESTIONS FROM THIS PAPER

1. Can T_train be quantified for an organization — say as the onboarding time until the first independent assessment of an out-of-unit proposal?
2. Is there a lower bound of legibility below which delegation is fundamentally unanswerable?
3. How many memory tiers does an organization need at minimum — do three suffice, or does continuity collapse without the fourth?

# NOTES

**On method.** Conception, argument and responsibility rest with the author. An AI system (Claude, Anthropic) was used as a working instrument for research, structuring, the transposition of the AGI-26 contributions and the bilingual drafting, with the explicit instruction to assess the relation of LEAD21 to AGI critically and with the greatest possible objectivity, including against the author's interest. Readers are expressly invited to evaluate the contribution of LEAD21 to AGI critically, with maximum objectivity of their own, and to draw their own conclusions.

**Declaration of interest.** The author is the developer and owner of LEAD21 and has a commercial interest in its dissemination. Every statement about LEAD21 is assigned a degree of coverage and an evidence status.

**On sources.** The author evaluated the posters he documented at AGI-26 as well as titles, abstracts and the proceedings tables of contents — not the full texts. Verbatim quotations come from the posters.

# REFERENCES


**AGI-26**

Botnick, M., Hammer, P., Isaev, P., Goertzel, B., & Crawford, K. (2026). OmegaClaw: A continually operating agentic architecture under bounded resources. In M. Iklé, A. Franz, A. Kemendo & H. Lowy (Eds.), *Artificial General Intelligence. AGI 2026, Part I* (LNCS 16854, pp. 338–353). Springer. https://doi.org/10.1007/978-3-032-33010-9_22

Buteau, T., & Kommrusch, S. (2026). Raman, not varelse — Why acceptably controllable AGI policies cluster around human-like stacks. In Iklé et al. (Eds.), *AGI 2026, Part I* (LNCS 16854, pp. 134–153). Springer.

Diamond, J. (2026). HetzerkVSM: Categorical vector-symbolic memory for typed knowledge graphs and simulation traces. In Iklé et al. (Eds.), *AGI 2026, Part I* (LNCS 16854, pp. 206–221). Springer.

Montes, G. A. (2026). Trace fields as externalized memory: A stigmergic multi-agent simulation of distributed control and trace-ecology governance. In Iklé et al. (Eds.), *AGI 2026, Part II* (LNCS 16855, pp. 78–90). Springer.

**Organizational research**

Nelson, R. R., & Winter, S. G. (1982). *An evolutionary theory of economic change.* Harvard University Press.

Walsh, J. P., & Ungson, G. R. (1991). Organizational memory. *Academy of Management Review, 16*(1), 57–91.

**External research**

Christensen, C. M., et al. (2011). The new M&A playbook. *Harvard Business Review*: 70–90% of M&A fail to deliver the expected value.

McKinsey (2024). *The culture compass*, survey of roughly 1,100 M&A leaders: 44% name cultural friction as the main reason.