Clarify operational readiness first. Then scale your agents.

When "Agentic AI Readiness Check" is especially valuable

Agents plan, access systems and carry out tasks on their own. Technical feasibility is therefore not yet operational readiness: what an agent may do, what it can access and who is accountable all have to be settled. Start without verified prerequisites and you fund pilots that never reach production.
Specifically, the aim is to …

DECIDE	BECAUSE …
• Whether your organization can run an agent in day-to-day operations	• the bottleneck sits between the agent and the organization, not in the model.
• Which gap has to be closed before you start	• the weakest capability decides the outcome.
• Which initiatives you deliberately defer	• without prioritization, budget drains into pilots.

VALUE LEVER 1 · DESIRED OUTCOME

An evidence-based readiness profile — not a maturity certificate.

You receive a basis for decision-making — not a snapshot of self-assessments:

- **Seven capabilities are assessed with evidence.** Strategy & Value Rationale · Processes & Workflows · Data & Knowledge · Technology & Architecture · Governance & Accountability · Skills, Roles & Organization · Economics & Scaling. 42 checkpoints against defined anchors from 1 to 5, backed by evidence rather than self-assessment.
- **Your bottleneck is named.** For every gap you receive the cause, the gate it affects and the action required — with an owner and a due date.

VALUE LEVER 2 · HIGH DECISION CONFIDENCE

Your verdict holds up, because a bottleneck cannot be averaged away.

A strong overall profile does not make an organization ready. That is why the Readiness Check applies three gates:

- **Secure completeness.** Every relevant checkpoint is assessed; "not relevant" is a documented scope decision, not a zero.
  Result: a verdict with no blind spots.
- **Test pilot readiness.** No relevant single checkpoint may fall below level 3 — neither an average nor a weighting compensates for it.
  Result: the blocker stays visible instead of disappearing into the mean.
- **Test scalability.** Every required capability must reach the target level that fits your scope, target autonomy and risk context.
  Result: a substantiated decision — STOP, PILOT or SCALE.

Copyright Strat-Exx 2026

VALUE LEVER 3 · FAST DECISION

In approx. three weeks, you have an evidence-based readiness verdict.

The check shortens the path from an open agentic ambition to a well-founded decision on whether to start.

- **Week 1: Set the reference frame.** Scope, planned use cases, target autonomy and risk context are defined.
- **Week 2: Gather evidence.** Six to ten interviews and a systems review; every rating is backed by evidence.
- **Week 3: Calibrate and decide.** Gates, bottlenecks and actions are made binding in the decision workshop.

> **Check principle.** First set the reference frame, then assess with evidence, then decide. Speed comes from a fixed assessment framework — not from less rigor.

VALUE LEVER 4 · LOW MANAGEMENT EFFORT

About half a day per interview partner — no more.

Strat-Exx runs the interviews, reviews the systems, calibrates the ratings and facilitates the decision — without workshop series or parallel projects.

Only four inputs are required from your organization:

- Access to process, system and governance documentation,
- interview partners from the business, IT, data management and compliance,
- existing evidence: process models, policies, logs, budgets, architecture diagrams,
- management decisions on scope, target level and the action plan.

VALUE LEVER 5 · LIMITED OPERATIONAL RISK

No action rights until control over them is proven.

The greater risk is not an early stop decision. It is granting an agent rights in live systems before oversight, traceability and reversibility have been demonstrated.

The check limits this risk:

- Every capability is assessed against the same anchors — no pet projects.
- A single checkpoint below level 3 stops pilot approval.
- Weak evidence makes the result explicitly provisional — it does not quietly change the score.
- STOP carries the same weight as SCALE.

Copyright Strat-Exx 2026

> **Methodological note.** The Readiness Check follows the logic of the NIST AI Risk Management Framework and the risk-based classification of the EU AI Act — designed to move the start decision from opinion to verifiable evidence.

By project end, all the evidence for decision-making is on the table:

A. a **readiness profile** across seven capabilities with 42 assessed checkpoints,
B. an **evidence profile** for each rating, from unsubstantiated to tested under real conditions,
C. a **bottleneck analysis** with cause, gate impact and priority for each capability,
D. a substantiated **management verdict**: STOP, PILOT or SCALE,
E. an **action plan** with target state, owners, expected evidence and due dates,
F. a **Route to Agentic AI** in four waves, each with a gate before the next.

> **The benchmark.** It is not the average that decides but the weakest capability in the overall interplay — and whether its gap has an owner, an action and a due date.

Which of the seven capabilities would slow your agentic initiative most today?

YOUR CONTACT



Dr. Clemens Oertel

CONFIDENTIAL STRATEGY CONVERSATION

30-MINUTE CONVERSATION

Via Google Meet · book anytime

calendly.com/clemens-oertel-strat-exx/30min

SCAN & BOOK

Copyright Strat-Exx 2026