



AI Readiness Thesis

Readiness is not a score; it is a shared view of the conditions that make a useful first move possible.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 50 / EMBER & MINT

EXECUTIVE WORKING EDITION - 08 OCT 2026

50

AT A GLANCE

Readiness is not a score; it is a shared view of the conditions that make a useful first move possible.

Readiness is often sold as a maturity score that ranks you against others. We think that framing is unhelpful. What matters is whether the conditions around a specific first move are good enough: how people already use AI, what skills and information they have, how work is organized, and what guardrails exist. This thesis describes four stages (See, Name, Strengthen, Review), a diagnostic across five dimensions, who should own the work, how to track change, and when readiness work is the wrong thing to do.

WHO IT IS FOR

Operations, technology, people and risk leaders who want to understand the conditions around AI before asking staff to change workflows or committing to a tool.

TIME TO USE

50 minutes to read; two to three sessions to apply

What you will have at the end

- A shared, plain-language view of strengths and gaps across five dimensions
- A short list of the smallest changes that make a first experiment safer and more useful
- A clear statement of what can start now and what should wait
- A living baseline you can revisit every quarter

Contents

01	Readiness is a view of conditions, not a grade	03
02	What most organizations get wrong about readiness	04
03	Four stages: See, Name, Strengthen, Review	05
04	See what is actually happening	07
05	Name the constraints and enabling conditions	08
06	Strengthen the smallest things that matter	10
07	Review: keep a living baseline	11
08	Diagnostics and ownership	12
09	Measuring progress without false precision	14
10	Risks, trade-offs and when not to do this	15
11	How Halden works with teams	16
12	A first quarter that produces evidence	17
13	Terms used in this publication	18

THE POINT OF VIEW

Readiness is a view of conditions, not a grade

01



The useful question is not how ready you are in general. It is whether the conditions around one specific first move are good enough to try it.

We hold that readiness is relative to a purpose. An organization can be well prepared to use an assistant for drafting internal summaries and poorly prepared to let one answer customer questions. A single score hides that difference. It also invites the wrong conversation, which is whether you are above or below average, instead of what you could safely do next week.

Readiness work is therefore an act of visibility. It makes current use, confidence, priorities, information and governance concerns visible to the people who must make decisions. Often the most valuable finding is not a gap but a surprise: AI is already being used more widely than leaders assumed, or the trusted source everyone relies on turns out to be one person's folder.

Readiness is also never finished. People learn, tools change, rules are updated and work is redesigned. A baseline taken today is useful because it can be compared with one taken in a quarter. We prefer a living baseline, light enough to refresh, to a large one-time audit that is out of date by the time it is presented.

- Readiness is judged against a specific first move, not in the abstract
- The goal is shared understanding, which matters more than the number
- Gaps are described in terms of the smallest change that closes them
- Strengths are named, because they are what you build on
- The baseline is revisited, so it stays alive

A USEFUL TEST

If your readiness work ends with a color-coded chart and no decision about what to start, stop, or fix first, it has described the situation without helping anyone act on it.



WORKING QUESTION

What is the specific first move you would like to be ready for, and who besides you would define it the same way?

COMMON FAILURES

What most organizations get wrong about readiness

02

Mistakes usually come from treating readiness as a gate, a technology question or a one-off audit.

The most common mistake is treating readiness as a gate that must be fully passed before anything can begin. This produces long preparation phases that delay learning. In reality many conditions improve fastest through contained use. Waiting for perfect data governance before any experiment usually means neither the experiment nor the governance happens.

The opposite mistake is ignoring readiness entirely and starting with a tool. Organizations then discover constraints in the middle of a pilot: documents nobody can find, staff who are unsure what they may paste into a tool, or a team with no time to change how they work. These are cheaper to find in a one-week readiness check than in the fourth week of a failing pilot.

A third mistake is asking only leaders. Executives tend to see strategy and policy. Practitioners see the tools actually in use, the real state of the documents and the unwritten rules. A readiness picture built from one level of the organization is reliably wrong in a predictable direction: it is more orderly than the truth.

FAILURE PATTERN	WHAT IT LOOKS LIKE	BETTER APPROACH
Readiness as a gate	Months of preparation before any use	Define the smallest safe first move and prepare for that
Readiness as a score	A rank against other firms	A description of conditions and next changes
Leaders-only view	Survey of the executive team	Include practitioners and managers

FAILURE PATTERN	WHAT IT LOOKS LIKE	BETTER APPROACH
Technology-only view	Focus on platforms and security alone	Cover people, information, process and decisions
One-off audit	A report that is never refreshed	A light baseline reviewed quarterly
Gaps without fixes	A long list of weaknesses	Each gap paired with a small, owned action

ILLUSTRATIVE EXAMPLE

A professional services firm (illustrative) assumes AI use is minimal. A short, anonymous check shows many staff use public tools weekly, with no guidance on client information. The urgent task is a simple usage note, not a platform.

WORKING QUESTION

Which view of your organization is missing from your current picture: leaders, managers or the people doing the work?

THE METHOD

Four stages: See, Name, Strengthen, Review

03

The method moves from observation to a small set of changes, and then back to observation so the picture stays current.

See gathers evidence about current use, confidence, priorities, workflow visibility and governance concerns. Name turns the evidence into explicit statements about constraints and enabling conditions across people, information, process and decisions. Strengthen selects the smallest changes that make a contained experiment safer and more useful. Review sets up the living baseline and the rhythm for revisiting it.

We organize the evidence around five dimensions, which are also the dimensions scored by our free readiness assessment at halden.company/assessments: current use, skills, workflows, information and governance. They are not a complete model of an organization. They are a practical set of questions that covers most of what decides whether a first move goes well.

STAGE	CORE QUESTION	MAIN OUTPUT
1. See	What is actually happening, and what do people think?	Evidence on use, confidence and concerns
2. Name	Which conditions help or hinder a first move?	A short statement of strengths and constraints
3. Strengthen	What are the smallest changes that matter?	A prioritized list of owned, small actions
4. Review	How will we know things have changed?	A living baseline and a review date

DIMENSION	WHAT IT COVERS	TYPICAL EVIDENCE
Use today	Which tools are used, how often, and what	Short survey, tool inventory, conversations
Skills	Ability to prompt, check output, share practice	Observed behavior, training history
Workflows	How AI use is chosen, measured and improved	Documented tasks, baselines, owners
Information	Findability, currency and permitted use of content	Source inventory, owner list
Governance	Guidance, accountability, response to errors	Policies, named owners, incident routes

WORKING QUESTION

Which of the five dimensions do you currently know least about, and who could tell you?

STAGE 1: SEE

See what is actually happening

04



Start with behavior and evidence, not opinions about maturity. Questions about what people did in the last 90 days are more reliable than questions about how good they are.

Seeing is mostly listening. The aim is a factual picture of who uses what, for which tasks, with what results, and how people feel about it. Self-rated skill is a weak signal because people differ in how generously they rate themselves. Ask about specific recent behavior, such as what happened the last time an answer was poor, and answers become more comparable.

01 Define the scope

Choose the teams or functions to cover and the first move you are trying to be ready for.

02 Collect use data

Run a short, anonymous survey and a few conversations about tools tried, frequency and tasks. Include an option for not sure.

03 Look at the work

Ask which workflows people think AI could help, and whether any current use has been measured against a starting point.

04 Inspect the information

List the five or six sources people rely on most and check who owns them and how current they are.

05 Read the rules

Gather existing guidance, approvals and any incident routes, written or informal.



Honest handling of uncertainty

Not sure is an important answer. If many people cannot say who owns a source or what guidance exists, that is a visibility gap and is itself a finding. Do not score uncertainty as failure, but do not hide it. Visibility gaps tend to be among the cheapest and most valuable things to fix.



COMMON MISTAKE

Running the survey and publishing results without saying what will happen next. Staff will answer less candidly next time. Say up front what you will do with the findings and who will see them.

WORKING QUESTION

If you asked ten people in different roles what AI they used last week, how many different answers would you expect?

STAGE 2: NAME

Name the constraints and enabling conditions

05

Naming turns scattered evidence into plain statements that leaders can discuss and act on.

Naming is the interpretive step. You read the evidence across the five dimensions and write short statements about what is helping and what is in the way. The most revealing statements come from patterns between dimensions, not from any single one. High use with weak governance is a different situation from low use with strong governance, even if both average to the same score.

PATTERN ACROSS DIMENSIONS	WHAT IT USUALLY MEANS	FIRST RESPONSE
Use ahead of governance	People are moving faster than the guardrails	A short, clear note on what may go where, plus a named owner
Skills ahead of workflows	Capability exists, but work has not been chosen	Pick a few workflows, give them owners and baselines
Workflows ahead of skills	Targets are identified but practice is thin	Pair each workflow with hands-on learning

PATTERN ACROSS DIMENSIONS	WHAT IT USUALLY MEANS	FIRST RESPONSE
Weak information	Answers will be poor and risk high	Tidy and assign owners to a handful of key sources
Strong governance, low use	Rules exist, but may feel like a sign	Add approved tools, examples and a quick route for questions

Write each statement in plain language and tie it to evidence. For example, many people use public tools weekly and none we asked knew of guidance on client information. Avoid vague statements such as culture is not ready. If you cannot attach a piece of evidence and a possible action, the statement is not yet useful.

- Name strengths first, and be specific about what to protect and build on
- Keep constraints to a short list, ideally no more than seven
- Distinguish blockers (must fix before a first move) from frictions (can be improved alongside)
- Mark what you do not know and how you could find out
- Check the statements with a few practitioners before sharing widely

WORKING QUESTION

Which single statement about your organization would most change a decision if it turned out to be true?

STAGE 3: STRENGTHEN

Strengthen the smallest things that matter

06

Choose the minimal changes that make a contained experiment safer and more useful. Leave the rest for later.

The temptation after a readiness exercise is to produce a long program of improvement. Resist it. Most organizations get most of the benefit from a handful of small, owned changes made quickly: a one-page usage note, an approved tool list, an owner for each of the top sources, a protected hour a week for a pilot team, and a route for questions.

01 Pick the first move

Confirm the contained experiment you want to enable, with its users, task and boundary.

02 List the blockers

From your named constraints, identify those that must change before the experiment can begin safely.

03 Choose small fixes

For each blocker, define the smallest action that would be enough, with an owner and a date.

04 Decide what waits

Record the gaps that can wait, so they are not forgotten or confused with blockers.

05 Communicate

Tell staff what is changing and why, in the same plain language used in the findings.

GAP	SMALLEST USEFUL CHANGE	LIKELY OWNER
No guidance on information handling	A one-page note on what may beRisk, legal or operations lead entered into which tools	
No accountable person	Name an owner for AI use and fo incidents	Executive sponsor
Unreliable sources	Assign owners and review dates for knowledge or data lead the top sources	

GAP	SMALLEST USEFUL CHANGE	LIKELY OWNER
No practice time	Protect a small, regular block for pilot team	Team manager
No measure	Capture a simple baseline before the experiment	Business owner

ILLUSTRATIVE EXAMPLE

A logistics company (illustrative) wants to try summarizing customer emails. The blockers are an unclear rule on customer data and no agreed definition of a good summary. Both are fixed in a week. The remaining gaps wait.

WORKING QUESTION

What is the one change you could make in the next five working days that would most reduce risk for people already using AI?

STAGE 4: REVIEW

Review: keep a living baseline

07

Readiness decays and improves. A light, repeated review keeps the picture honest and shows whether the changes worked.

A baseline is only useful if it can be compared with something. The review stage defines what you will measure again, how often, and who looks at it. Quarterly works for most organizations. The same questions asked in the same way allow you to see movement, and the changes in answers will often be more informative than the absolute level.

Keep the review light. The goal is to maintain a shared view, not to run a program. Re-ask the core questions, check the status of each action from the Strengthen stage, and add new evidence from pilots and training. Where readiness has improved enough to support a larger move, say so explicitly and update the first-move definition.

Review checklist

- The same core questions are asked in the same way each cycle
- Every Strengthen action has a status and an owner
- Changes since the last review are written down with evidence
- New constraints from pilots or incidents have been added
- Findings are shared with the people who were asked
- The next first move, or the next larger step, is defined

COMMON MISTAKE

Treating an improved score as success. Check that behavior changed, such as fewer unsanctioned uses or faster answers to staff questions, not only that people rated themselves higher.

WORKING QUESTION

Who will own the next review, and what date is it in the calendar right now?

DIAGNOSTICS AND ROLES

Diagnostics and ownership

08

A short self-check surfaces the dimension that needs attention first. Named roles keep readiness from becoming an orphan task.

Use the checklist below with a mixed group of leaders, managers and practitioners. Answer yes, partly, no or not sure. Then compare answers by role, since gaps between how leaders and practitioners respond are among the most useful findings. Our free readiness assessment at halden.company/assessments turns these dimensions into a structured result with insights from the pattern across areas.

Readiness self-check

- We know which AI tools are used in the organization and by roughly how many people
- People can describe what they do when an AI answer is poor
- Some structured learning on using AI at work has taken place
- People share what works with one another
- We chose where to use AI deliberately, and at least one use has a baseline
- Someone notices and improves the process when an AI-assisted task goes wrong
- People can find current, trustworthy information on our products, policies and processes
- People know what information may go into an AI tool
- There is written guidance, and a named person accountable for AI use
- We know what would happen if an AI tool produced a harmful error tomorrow



ROLE	RESPONSIBILITY	TYPICAL HOLDER
Readiness lead	Runs the four stages and maintains the baseline	Operations, people or transformation lead
Executive sponsor	Backs the findings and the actions that follow	A member of the leadership team
Information owner	Maintains key sources and clarifies permitted use	Knowledge, data or IT lead
Risk contact	Provides guidance and an escalation route	Legal, compliance or security lead
Team representatives	Bring practitioner reality into the picture	Managers and respected practitioners



The readiness lead does not need to be technical. They need credibility with practitioners, access to leaders and the patience to do a short survey properly. In smaller organizations this is often the operations or people manager, with a technical adviser consulted as needed.

WORKING QUESTION

Where would your answers and your front-line staff's answers differ most on this checklist?



MEASURING PROGRESS

Measuring progress without false precision

09

Track a few indicators of real change and be honest about what they can and cannot tell you.

Good readiness indicators describe conditions, not enthusiasm. Examples include the share of staff who know where guidance is, the number of key sources with a named owner and a review date, the time between a question about AI use and an answer, and the number of use cases with a baseline. Each can be counted without elaborate tools and each ties to an action you took.

INDICATOR	WHAT IT SHOWS	HOW TO COLLECT
Staff who can find the guidance	Whether governance is usable	Two-question pulse each cycle
Key sources with owned information and review date	Information reliability	Source register
Use cases with a baseline	Evidence discipline	Opportunity register
Time to answer a use question	Support responsiveness	Log of questions and replies
Reported uncertainty (not sure answers)	Visibility	Compare each cycle's survey
Actions completed on time	Follow-through	Action tracker

Interpret movement carefully. A rise in reported tool use after you publish guidance may reflect people feeling safe to say what they do, not a real jump in use. A fall in confidence after training can mean people now understand the limits. Pair numbers with a few short conversations so you can read them correctly.

Avoid targets that reward gaming, such as a required score by year end. A readiness measure that people are pressured to improve will improve, whether or not the underlying conditions do. Report trends and actions, and discuss them as a group.

WORKING QUESTION

Which indicator would you trust enough to change a decision, and how would you collect it cheaply?



RISKS AND TRADE-OFFS

Risks, trade-offs and when not to do this

10



Readiness work has real costs in time, attention and expectation. Use it where it pays back.

The central trade-off is thoroughness against momentum. A more complete picture takes longer and risks delaying useful learning. A thinner picture risks missing a constraint. Our default is to size the work to the first move: if the move is low risk and contained, a few conversations and a short survey are enough. If it involves sensitive information or customers, go deeper before starting.

- Survey fatigue and distrust if findings are not shared or acted on
- Raising expectations that every gap will be closed quickly
- Anonymity limits: small teams can be identified, so protect responses carefully
- Treating a self-assessment as an audit, which it is not
- Legal, regulatory and contractual duties vary by sector and region: check current requirements and seek qualified advice, because a readiness exercise does not establish compliance



WHEN NOT TO DO A READINESS EXERCISE

Skip it if you already have a clear, contained experiment with low risk and a willing team: run the experiment and learn from it. Skip it if leaders have not yet agreed on the business question: that is strategy work and should come first.



Readiness work can also become a way of postponing. If three cycles have passed and no first move has started, the issue is probably decision-making or ownership, not readiness. Name that plainly and act on it.

**WORKING QUESTION**

What would you be tempted to keep preparing for, when a small, safe test could teach you more?

**WORKING WITH HALDEN**

How Halden works with teams

11

We help organizations see their position clearly and pick the next small step, and we say so when you can do it yourselves.



Halden & Company is an AI advisory and implementation partner. For readiness, we design and run the See and Name stages, offering anonymous collection, a neutral voice and experience reading patterns across the five dimensions. We then help you choose the Strengthen actions and set up the review cycle so your own team can maintain the baseline afterward.

- Free assessments at halden.company/assessments give you a structured starting picture in minutes
- Facilitated sessions can bring leaders and practitioners into one conversation
- We help write the plain-language findings and the short list of owned actions
- We will say if readiness work is not the right next step for you
- Your team owns the baseline and the method, with no dependency on us to refresh it



If you would like to discuss a specific situation, you can contact us at halden.company/contact. We would rather give you a clear, honest view of what to do first than sell you a larger engagement than you need.

WORKING QUESTION

What would make a readiness conversation with an outside party worth having for your team?



30 / 60 / 90

A first quarter that produces evidence

12



A first quarter that ends with a shared picture, a handful of fixed blockers and a baseline you will revisit.

First 30 days: ground it

- Define the first move you want to be ready for and its boundary
- Run a short anonymous check on use, skills, workflows, information and governance
- Hold conversations with practitioners in at least three roles
- Publish a one-page usage note if none exists

**Days 31 to 60: test it**

- Write the named strengths and constraints, separating blockers from frictions
- Choose small fixes for each blocker with owners and dates
- Assign owners and review dates to the top information sources
- Share findings and actions with everyone who was asked

Days 61 to 90: decide

- Start the contained first move with a baseline
- Re-ask the core questions and compare with the first baseline
- Review action status and add constraints found in practice
- Set the next quarterly review and name its owner



GLOSSARY

Terms used in this publication

13

TERM	MEANING
Readiness	The degree to which current conditions support a specific, contained first use of AI.
Baseline	A recorded picture of current conditions or performance, used for later comparison.
Living baseline	A baseline that is lightly refreshed on a schedule rather than produced once.
Dimension	One of five areas of evidence: use today, skills, workflows, information and governance.
Visibility gap	An area where people cannot say what is happening or who owns it, treated as a finding rather than a failure.
Blocker	A condition that must change before a contained first move can begin safely.
Friction	A condition that makes work harder but can be improved alongside a first move.
Contained experiment	A small trial with defined users, task, boundary and outcome.
Governance	The guidance, accountability and routes for decisions and incidents around AI use.
Information owner	A named person responsible for the accuracy and currency of a source.
Usage note	A short, plain document saying what information may be entered into which tools.

A note on use

This publication is a public orientation resource, not legal, technical, financial or professional advice. Adapt it to your people, information, workflows and obligations, and check current tool terms, regulations and risks before acting.

TAKE THE NEXT STEP

Take the free Halden assessment at halden.company/assessments, or talk to us at halden.company/contact about turning this into a plan for your organization.

