


State of the art

Agile Leadership – Decision Architecture

Decision-making in agile teams

Reference map of current professional knowledge



Updated June 2026

About this document

This document is a curated map of the knowledge, practices and models in professional use as of its publication date for designing how decisions are made in agile teams: where they are made, with what information, at what speed, with what safety net and what part AI plays at each point of the flow. For each entry, the map places it in context, indicates its degree of adoption in current professional practice and points to where to find reference information.

Use it as an observatory and reference point to check whether the professional knowledge you rely on is aligned with current practice and at the leading edge.

It is complemented by an in-depth companion guide that develops each concept and by the training and assessment platform in Skill Arena. In the [Agile Leadership – Decision Architecture Decision-making in agile teams](#) area you can test your level of knowledge and, if you pass, obtain a diploma that provides curricular accreditation of professional competence and currency in this area.



State of knowledge

Knowledge in this area evolves at two speeds. Its foundations —intent-based leadership, reversibility, decision rights, deciding under uncertainty— have been settled for decades and change slowly. The human–AI decision layer, by contrast, evolves extremely fast: frameworks, adoption data and regulation move month by month, and account for most of the revisions of this map.

Throughout the document, the labels (ESTABLISHED, CONSOLIDATING, EMERGING) help identify the maturity of each concept:


ESTABLISHED settled consensus; knowledge taken as required.

CONSOLIDATING gaining adoption fast; not yet universal but already relevant.

EMERGING recent frontier; high relevance and high volatility.

Block 1 — From the leader who decides to the leader who designs

The difference between teams that work and teams that get stuck lies not in the leader's charisma, but in how decisions are designed and distributed.



From the apex that concentrates decisions to the network that distributes them by design.

Decision architecture as the object of agile leadership · CONSOLIDATING

Traditional leadership consists of making decisions; agile leadership consists of designing the system that makes them: where they are made, with what information, at what speed and with what safety net. A decision is a flow of the system, not an attribute of the boss: every hierarchical hop adds latency and loses context. The uncomfortable truth is that agile teams almost never have an explicit decision map; the architecture always exists, but it has grown on its own, like technical debt.

Why it is here now. Generic inspirational discourse (vision, purpose, servant leadership) is true but inoperative: it does not say what to do on Monday. The 2026 trend reports single out "decision velocity" —deciding with incomplete information— as the number one competence of the year (DHR Global), which shifts the focus from charisma to the design of the decision system.

Where to look. [DHR Global — Talent Trends 2026](#) · [Agile leadership: don't make the decisions, design how they are made — Scrum Manager](#)

Decision Intelligence: the emerging discipline · EMERGING


Decision Intelligence is the discipline of turning information into better decisions, combining data science with the social and management sciences. The term, coined by Cassie Kozyrkov during her time at Google, today coexists with two senses that are worth keeping apart: the professional leadership competence this topic develops, and the category of technology platforms (modelling, governance and execution of decisions with AI at enterprise scale) that Gartner established as a market with its inaugural Magic Quadrant of January 2026.

Why it is here now. The consolidation of the term as a market category makes it necessary to anchor its competence sense: DI platforms will reach organisations and the agile leader will need judgement to decide which decisions are delegated to them, with the same logic applied to the rest of the decision map. Gartner already describes the market as late emerging, in accelerating adoption.

Where to look. [Kozyrkov — Introduction to Decision Intelligence](#) · [Gartner — MQ Decision Intelligence Platforms \(2026\)](#)

Block 2 — Moving authority to the information

Instead of moving information towards authority, move authority to where the information is.



Authority travels towards the information, instead of pushing information up towards authority.

Intent-based leadership · ESTABLISHED

David Marquet's principle, forged in command of the submarine Santa Fe and described in *Turn the Ship Around!*: instead of asking permission or waiting for orders, whoever holds the information states "I intend to..." and authority is transferred progressively and under control. It is the real operating basis of "self-organising teams": self-organisation is not the absence of authority, but authority placed where the information is.

Why it is here now. It is the conceptual foundation of the whole map and one of the most widely adopted leadership references of the past decade. Its concrete mechanics (gradual transfer through the language of intent) remains the most operative translation of the principle into everyday agile work.

Where to look. [Marquet — Turn the Ship Around!](#) · [Intent-Based Leadership resources](#)

The two pillars: competence and clarity of intent · ESTABLISHED

Decentralising is not handing out decisions and wishing everyone luck. Pillar 1: technical competence (whoever decides has to know how). Pillar 2: clarity of intent (they have to know what for). Without both, it is not decentralising: it is abandoning. The military "commander's intent" supplies the format: what is sought, what constraints apply and what is non-negotiable; the discipline consists of replacing instructions with intent.

Why it is here now. It is the same discipline that writing a good specification for an AI demands today: a spec is intent formalised for an agent, just as commander's intent is for a team. This connection turns a classic delegation practice into a dual-use competence in teams working with AI (see the topic SDD in agile teams).

Where to look. [SDD in agile teams — Skill Arena](#) · [Marquet — Turn the Ship Around!](#)

The team's decision map · CONSOLIDATING

The decision inventory as a starting practice: which decisions are made recurrently, who makes them today and who has the best information to make them. It works well as a single-topic retrospective.


The symptoms of an implicit architecture are recognisable: answers such as "it depends", "ask so-and-so" or silence in the face of questions like who can change a library, drop a story or delay a delivery.

Why it is here now. It is the practice that turns the thesis of block 1 into something actionable: you cannot redesign an architecture that has not been made visible. It is also the artefact later extended to hybrid human–AI teams (block 8), so its adoption gains relevance as agents arrive.

Where to look. [Management 3.0 — Delegation Poker](#) · [Agile leadership: don't make the decisions, design how they are made — Scrum Manager Observatory](#)

Block 3 — Reversibility as a design criterion

Classify by reversibility, not by importance: "important" is subjective; reversible or irreversible is operative.



One-way doors —crossed slowly— and two-way doors, which allow you to go back.

Reversible and irreversible decisions · ESTABLISHED

Jeff Bezos's distinction in his letters to shareholders: two-way doors (decisions that can be undone) as against one-way doors (irreversible or nearly so). The two mistakes are symmetrical: treating every decision as irreversible (committees, approvals, paralysis over things that can be undone in an afternoon) and treating the irreversible ones lightly (the monolith impossible to migrate, the contract with no way out). The design rule: reversible decisions go down to the level closest to the information, with permission to be wrong; irreversible ones go up only as far as needed, are deliberately slowed down and are documented.

Why it is here now. It is the most operative classification criterion of the map and the one that articulates the rest: it decides which mechanism to use (block 4), when to apply the 70% heuristic (block 5) and what can be delegated to an agent (block 8). Its original formulation remains the reference.

Where to look. [Bezos — 2015 letter to shareholders](#) · [Bezos — 2016 letter to shareholders](#)

Turning irreversible decisions into reversible ones · ESTABLISHED


The boundary between reversible and irreversible is not fixed: a good part of engineering work — tests, feature flags, incremental deployments, decoupled architectures— consists of making undoing cheaper. Every time undoing becomes cheaper, the zone in which the team decides without asking permission widens. The key conclusion: technical infrastructure is also decision architecture. The principle travels beyond software: exit clauses, pilots, phased commitments.

Why it is here now. It reframes common technical investments as organisational design decisions, which gives them a business case many teams do not know how to articulate. It is the constructive side of reversibility: not only classifying decisions, but moving decisions from one category to another.

Where to look. [One-way door decisions in software development](#) · [Bezos — 2015 letter to shareholders](#)

Block 4 — Decision-rights frameworks and record-keeping

Formalise recurrent or contentious decisions; do not bureaucratised the trivial ones.



Explicit decision rights per decision and per role, with their record: what was decided, what was ruled out and what was known.

Frameworks for assigning decision rights · ESTABLISHED

DACI (Driver, Approver, Contributors, Informed) and RAPID (Recommend, Agree, Perform, Input, Decide) as reference frameworks for making explicit who plays which part in each decision: DACI for cross-cutting day-to-day decisions, RAPID for complex multi-level governance. Delegation poker and the seven levels of delegation (Management 3.0) serve as a tool for explicit calibration between leader and team.

Why it is here now. They are the frameworks with the widest adoption and the most stable documentation in the area, and the practical answer to the "nobody knows who decides" symptom of pointer 2.3. The criterion for using them — formalise what is recurrent or contentious, not what is trivial— avoids the antipattern of bureaucratising agility.

Where to look. [Atlassian — DACI](#) · [Bain — RAPID](#) · [Management 3.0 — Delegation Poker](#)

Consensus, consent and the advice process · CONSOLIDATING

Alternatives to the boss-decides / team-votes pairing: sociocratic consent (things move forward when nobody has a reasoned objection) as against consensus (everybody agrees), and the advice process (whoever spots the problem decides, after consulting those affected and the experts). Each mechanism has a different cost and a different speed; the craft lies in matching mechanism to type of decision, using the reversibility of block 3 as the criterion.

Why it is here now. Teams that decentralise soon discover that "deciding together" without an explicit mechanism degenerates into permanent consensus, the most expensive and slowest of all. These mechanisms, coming from sociocracy and from self-managed organisations, are entering the standard repertoire of agile teams.

Where to look. [Sociocracy For All — consent decision-making](#) · [Comparing decision-making methods \(advice process\)](#)

Decision records: ADRs and lightweight formats · ESTABLISHED

Michael Nygard's Architecture Decision Records as a well-known example in software of a minimum viable format for any field: what was decided, what alternatives were ruled out and what was known at the time. It is not bureaucracy: it is what makes it possible, six months later, to tell a bad decision from a reasonable decision with bad luck. Without a record, every review of decisions is a verdict delivered with Monday's newspaper in hand.

Why it is here now. The original 2011 format remains valid and its adoption has consolidated beyond software architecture as a lightweight record of decisions of any kind. It is also a precondition of pointer 4.4: without a record there can be no decision retrospective.

Where to look. [Nygard — Documenting Architecture Decisions](#) · [adr.github.io](#)

Decision retrospectives · CONSOLIDATING


Retrospectives review what we did and how we coordinated, but rarely how we decided. The practice: take two or three decisions from the last quarter and evaluate the process, not the outcome. Did we have the information it was reasonable to have? Did the right person decide? Did we take the appropriate amount of time? Hindsight bias and outcome bias are the traps to neutralise: a good outcome does not validate a bad process, nor the other way round.

Why it is here now. It closes the inspect-and-adapt loop on the decision system itself, which is exactly what agility applies to everything else and almost never to itself. It gains relevance as judgement becomes the bottleneck (block 6): what is scarce has to be trained deliberately.

Where to look. [Fowler — Architecture Decision Record](#)

Block 5 — Deciding under uncertainty

If you wait until you have all the information, you are not deciding: you are confirming.



The OODA loop: the decisive phase is not Decide but Orient — the filter through which what is observed is interpreted.

The cost of waiting versus the cost of being wrong · ESTABLISHED

The useful question is not "do we have enough data?" but "does waiting cost more, or does being wrong cost more?". The 70% heuristic (Bezos): with 70% of the information you would like, decide; the remaining 30% is bought more cheaply by acting than by analysing. Its applicability is conditional on the reversibility of block 3: for reversible decisions, waiting almost always costs more.

Why it is here now. It gives operative shape to the "decision velocity" that the trend reports place as the central competence of 2026: decision speed is not haste, it is having calculated the cost of waiting. It is also the antidote to the analysis paralysis that the abundance of data has aggravated, not solved.

Where to look. [Bezos — 2016 letter to shareholders](#) · [DHR Global — Talent Trends 2026](#)

The OODA loop and the primacy of Orient · ESTABLISHED

John Boyd's loop — Observe, Orient, Decide, Act — as a description of agile work: the winner is not whoever decides better in the abstract, but whoever iterates the complete loop faster. The nuance usually lost: the decisive phase is not Decide but Orient, the filter of experience, context, culture and mental models through which what is observed is interpreted. Two teams look at the same dashboard and see different things. Going very fast with the wrong orientation is not agility.

Why it is here now. Boyd published no book: the loop is known through his briefings and through canonical secondary sources, which explains why it often circulates simplified down to "speed". Recovering the primacy of Orient is the conceptual groundwork for block 6: it is precisely the phase AI does not solve.

Where to look. [OODA loop \(overview\)](#) · [The Tao of Boyd](#)

Cynefin as a decision contextualiser · ESTABLISHED

Dave Snowden's framework for matching type of context — clear, complicated, complex, chaotic — with mode of decision: good practice, expert analysis, probe-sense-respond experimentation or immediate action. Its central practical use: avoiding the mistake of applying complicated-domain


analysis to complex-domain problems. In the agile cadence, the sprint works as a probe in the complex domain.

Why it is here now. It remains the reference framework for answering the question that precedes any decision: what kind of context am I deciding in? Its canonical HBR article is still the standard way in, and its currency grows as it serves as the criterion for deciding which decisions admit automation and which demand experimentation.

Where to look. [Snowden and Boone — A Leader's Framework for Decision Making](#)

Block 6 — AI compresses the loop

AI makes Observe and Act cheaper; the bottleneck shifts to exactly the two middle phases: to judgement.



Asymmetric compression: observing and acting widen with AI; orienting and deciding become the bottleneck.

Asymmetric compression of the OODA loop · CONSOLIDATING

Generative AI makes Observe cheaper (summarising, analysing, spotting patterns in volumes previously out of reach) and Act cheaper (generating code, a document or a prototype in minutes). The loop turns faster than ever and the bottleneck shifts to Orient and Decide: to judgement. Validating faster than things are generated is the new constraint of the value chain.

Why it is here now. It is the structural change that justifies the AI layer of this map. The adoption data confirm it: according to Deloitte's 2026 human capital trends study, 60% of executives already use AI regularly to support their decisions, and Gartner projects that by 2027 half of business decisions will be augmented or automated by agents.

Where to look. [Deloitte — AI and the future of human decision-making](#)

The new economics of leadership · CONSOLIDATING

When producing was expensive, value lay in executing well what had been decided; when producing is cheap, value lies in deciding well what to produce. Whoever decides becomes the limiting factor of the value chain —in teams working with AI, the Product Owner evolves towards the AI Product Manager as the new strategic bottleneck. Decision saturation also appears: the problem is no longer a lack of information but the volume of competing decisions; filtering and assigning ownership becomes a structural function of leadership.

Why it is here now. "Human + AI" leadership is the central trend of 2026 according to DDI: questioning AI outputs, identifying biases and combining machine efficiency with human judgement ("AI fluency"). The evolution of roles connects with the guide Scrum in teams with AI, which develops the shift of the bottleneck towards judgement.

Where to look. [DDI — Leadership Trends 2026](#) · [Scrum in teams with AI — Skill Arena](#)

AI to open options, not to close them · CONSOLIDATING


Using AI as a generator of alternatives and as a devil's advocate — "what are we assuming?", "how could this go wrong?", premortem — as against using it as an oracle that hands over the final answer, which is a machine for manufacturing agreement. The practical rule: AI widens the menu; the team chooses the dish.

Why it is here now. It is the difference in use that separates AI which improves judgement from AI which replaces it without saying so. It takes up and extends the premortem competence of the skill AI applied to agile work, and it is the first-line countermeasure against the automation bias that opens block 7.

Where to look. [Deloitte — AI and the future of human decision-making](#)

Block 7 — Human judgement and oversight

Approving the AI's proposal without having been able to evaluate it feels exactly the same as deciding. But it is not deciding: it is signing off.



Real oversight: time and competence to review, and deliberate friction before irreversible decisions.

Automation bias: agreement is not judgement · CONSOLIDATING

When a system is right 90% of the time, the natural thing is to stop reviewing; the remaining 10% goes through without anyone looking at it. For human oversight to be real and not ritual, three conditions are needed: time, competence and the ability to dissent at no cost. If any is missing, the review is theatre.

Why it is here now. It is the first-line risk in teams adopting AI to decide, and the human face of the antipattern that Harness Engineering addresses in its sensors: exclusive faith in the reviewer. The conditions for real oversight are a matter of organisational design —of decision architecture—, not of individual goodwill.

Where to look. [Harness Engineering — Skill Arena](#) · [DDI — Hot Leadership Topics 2026](#)

Accountability in AI-assisted decisions · EMERGING

Who is accountable for a decision suggested by AI? The easy answer — "the human who approved it" — becomes a legal fiction when that human approves forty suggestions a day because the flow leaves no room for more. Responsibility without a real capacity to review is not responsibility: it is a lightning rod. Designing accountability requires sizing the volume of decisions to the real capacity for review and making explicit which decisions require substantive human evaluation.

Why it is here now. The regulatory framework for AI is converging on the principle of meaningful human oversight, and the 2026 agent authorisation frameworks (such as the World Economic Forum's ACAP) already formalise auditable and enforceable delegation. The distinction between real responsibility and a lightning rod moves from ethical debate to design requirement.

Where to look. [WEF — AI Agents in Action \(ACAP\)](#)

Deliberate friction as a safety mechanism · EMERGING

Removing friction from processes is generally a good idea; but a certain amount of friction in irreversible decisions is a safety mechanism. If AI removes it, it has to be reintroduced by design: not


slowness, but composure. Patterns: mandatory pauses, a second human reviewer, reflection periods for one-way decisions, "rate limiters" in automated flows. Managing the cadence of judgement is a new part of the leadership craft.

Why it is here now. The compression of the loop (6.1) also drags in decisions that should not be accelerated: speed stops being an achievement and becomes a design parameter. The principle ties in with the classic practice of deliberately slowing down one-way doors (3.1), now applied to flows where AI sets the pace by default.

Where to look. [Bezos — irreversible decisions and deliberate slowness](#)

Block 8 — Designing the hybrid human–AI architecture

What the agent decides, what it proposes for approval and what is exclusively human: the decision map, extended.



Three lanes of the hybrid map: the agent decides, propose-and-approve, and protected decisions that are exclusively human.

Delegating decisions to agents · EMERGING

Extending the team's decision map (2.3) to AI agents: which decisions an agent can make, which work in propose-and-approve mode and which are exclusively human. The criteria of the previous blocks apply unchanged: reversibility (3), demonstrated competence and clarity of intent (2), explicit decision rights (4).

Why it is here now. Frameworks for delegating to agents matured in 2026: proposals such as Google DeepMind's formalise the transfer of authority, responsibility and trust mechanisms in human–agent networks, and the WEF playbook contributes auditable authorisation profiles. The harness's guides (feedforward) and sensors (feedback) are the technical implementation of these decision rights: a direct connection with Harness Engineering.

Where to look. [DeepMind — Intelligent AI Delegation](#) · [WEF — AI Agents in Action \(ACAP\)](#) · [Harness Engineering — Skill Arena](#)

The risk of cognitive re-centralisation · EMERGING

The paradox: decisions are decentralised towards the teams... and the teams delegate them to one and the same AI model, with the same biases and the same "orientation". Decisions formally distributed, cognitively homogeneous: diversity of judgement was a property of the system that is not being protected. Countermeasures: diversity of sources and models, structured dissent (human red teams) and making the model's "orientation" explicit as one more input to be challenged.

Why it is here now. It is the least visible systemic risk of AI adoption: it shows up in no particular incident, but in the slow convergence of every decision towards one and the same filter. Whoever designed decentralisation to gain diversity of judgement discovers it can be lost through the back door.

Where to look. [DDI — Hot Leadership Topics 2026](#)

Preserving and training judgement · EMERGING

Judgement is trained by deciding small things and living with the consequences. If AI absorbs the small decisions, where will the judgement for the big ones come from? The risk is acute in junior profiles: AI cuts away the lower rungs of the ladder (minor decisions, cheap mistakes, growing judgement). The practice: protect small decisions so that people make them, even though AI could make them better; this is not inefficiency, it is the training cost of judgement —like hand-flying with the autopilot available.

Why it is here now. The 2026 analyses agree that, with scale and speed solved by AI, the constraint is human judgement: cultivating it deliberately becomes a design decision, with decision pathways for people development. It is the map's long-term investment: without it, blocks 6 and 7 are left without competent supervisors.

Where to look. [IMD — 2026 AI trends for leaders](#)

What to watch in the next revision

Signals of change the editorial team anticipates for the next edition of this map:

- The evolution of the Decision Intelligence platform market after Gartner's inaugural Magic Quadrant (January 2026): which types of decision start to be delegated to platforms and with what results, and how this affects the boundary between the competence sense and the product sense (pointer 1.2).
- The standardisation of frameworks for delegating to agents: academic proposals (DeepMind), authorisation profiles (the WEF's ACAP) and agent identity and authorisation initiatives in standards bodies. If a dominant framework converges, block 8 will have to reflect it.
- The regulatory development of the principle of meaningful human oversight in AI-assisted decisions, which may turn the design of accountability (7.2) from good practice into a requirement.
- The adoption data for assisted and automated decision-making: Gartner's projection (half of business decisions augmented or automated by agents in 2027) will determine whether the pointers of block 6 move from "consolidating" to "established".
- The consolidation of "decision velocity" and "decision saturation" in the 2027 trend reports, and the appearance of instruments for assessing decision competence in talent frameworks.

A note on the references

The links included were available and verified as of the update date. Given the pace of change in this area, some may change; each revision of the map also updates its references.

© 2026 Scrum Manager®. This work is published under a Creative Commons Attribution-NonCommercial 4.0 International licence (CC BY-NC 4.0). Official Scrum Manager trainers and centres are licensed under the terms of CC BY 4.0 for their training activity.