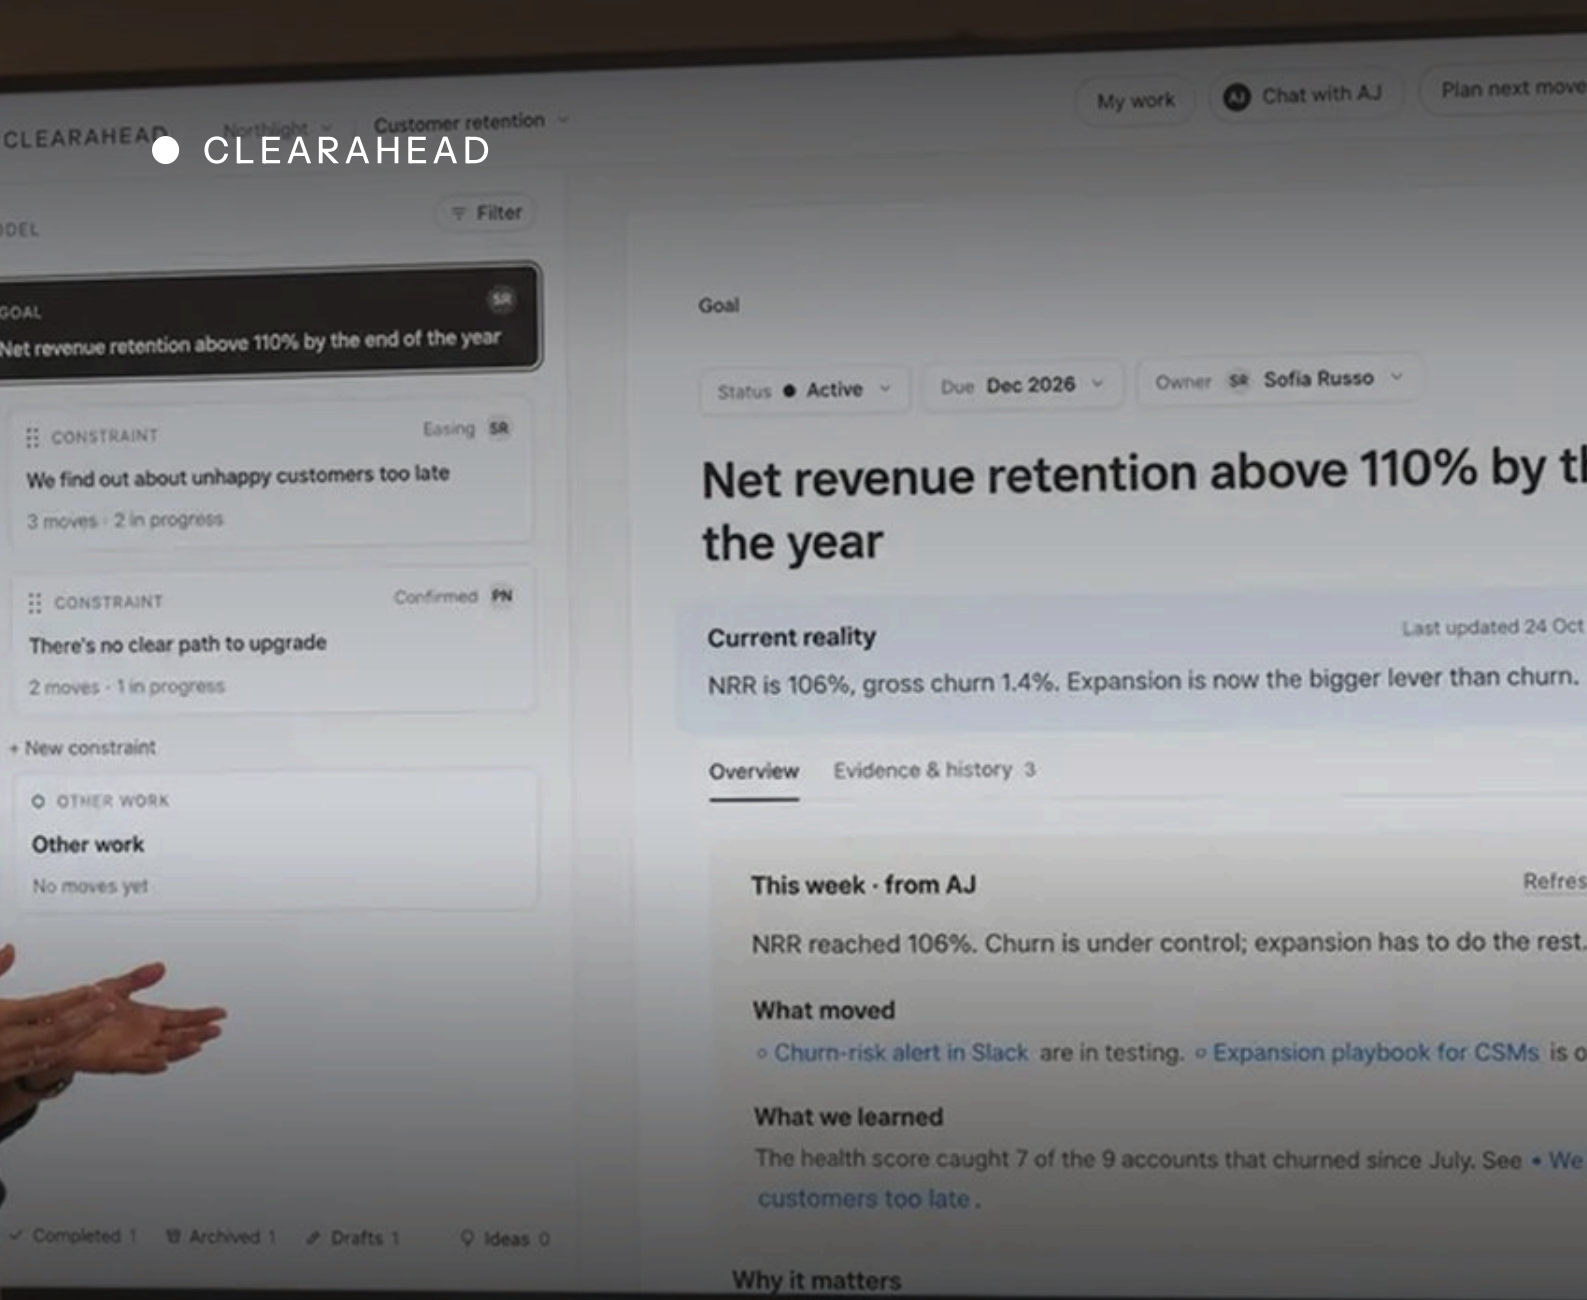
CLEARAHEAD



The plan is the bottleneck.

Why long-term strategy fails AI-enabled companies, and the cadence that replaces it.

A Clearahead white paper · Alex Durston, founder · November 2026

The short version

AI has made companies faster at almost everything except deciding what to do next, so for a growing number of businesses the slow, once-a-year machinery of long-term strategy has quietly become the constraint that holds the whole company back.

Strategic planning was already struggling long before anyone had used a large language model: most executives didn't believe it was worth the effort, most big decisions were made outside it, and most of the people executing a strategy couldn't say what it was. AI has widened that gap into a gulf, because the work now moves in days and weeks while the plan still moves in quarters and years.

This paper sets out the evidence, explains why speeding up execution does nothing for a company whose real constraint is deciding and aligning, and describes the alternative we've built Clearahead around: keep the long-term goal fixed as the target, let the strategy unfold as you clear whatever is holding that goal back, and run the whole thing on a cadence that matches the speed of the work.

11%

of senior executives strongly believe strategic planning is worth the effort.

Mankins & Steele, HBR, 2006

~7

months for the length of task AI can complete to double.

METR, 2025

95%

of organisations see no measurable P&L return from generative AI.

MIT Project NANDA, 2025

Contents

01 Faster at everything except deciding

02 A plan that was out of date before it started

03 Long-range planning was failing before AI arrived

04 What AI changed

05 The bottleneck moved

06 Why leaders hold on to the Gantt chart

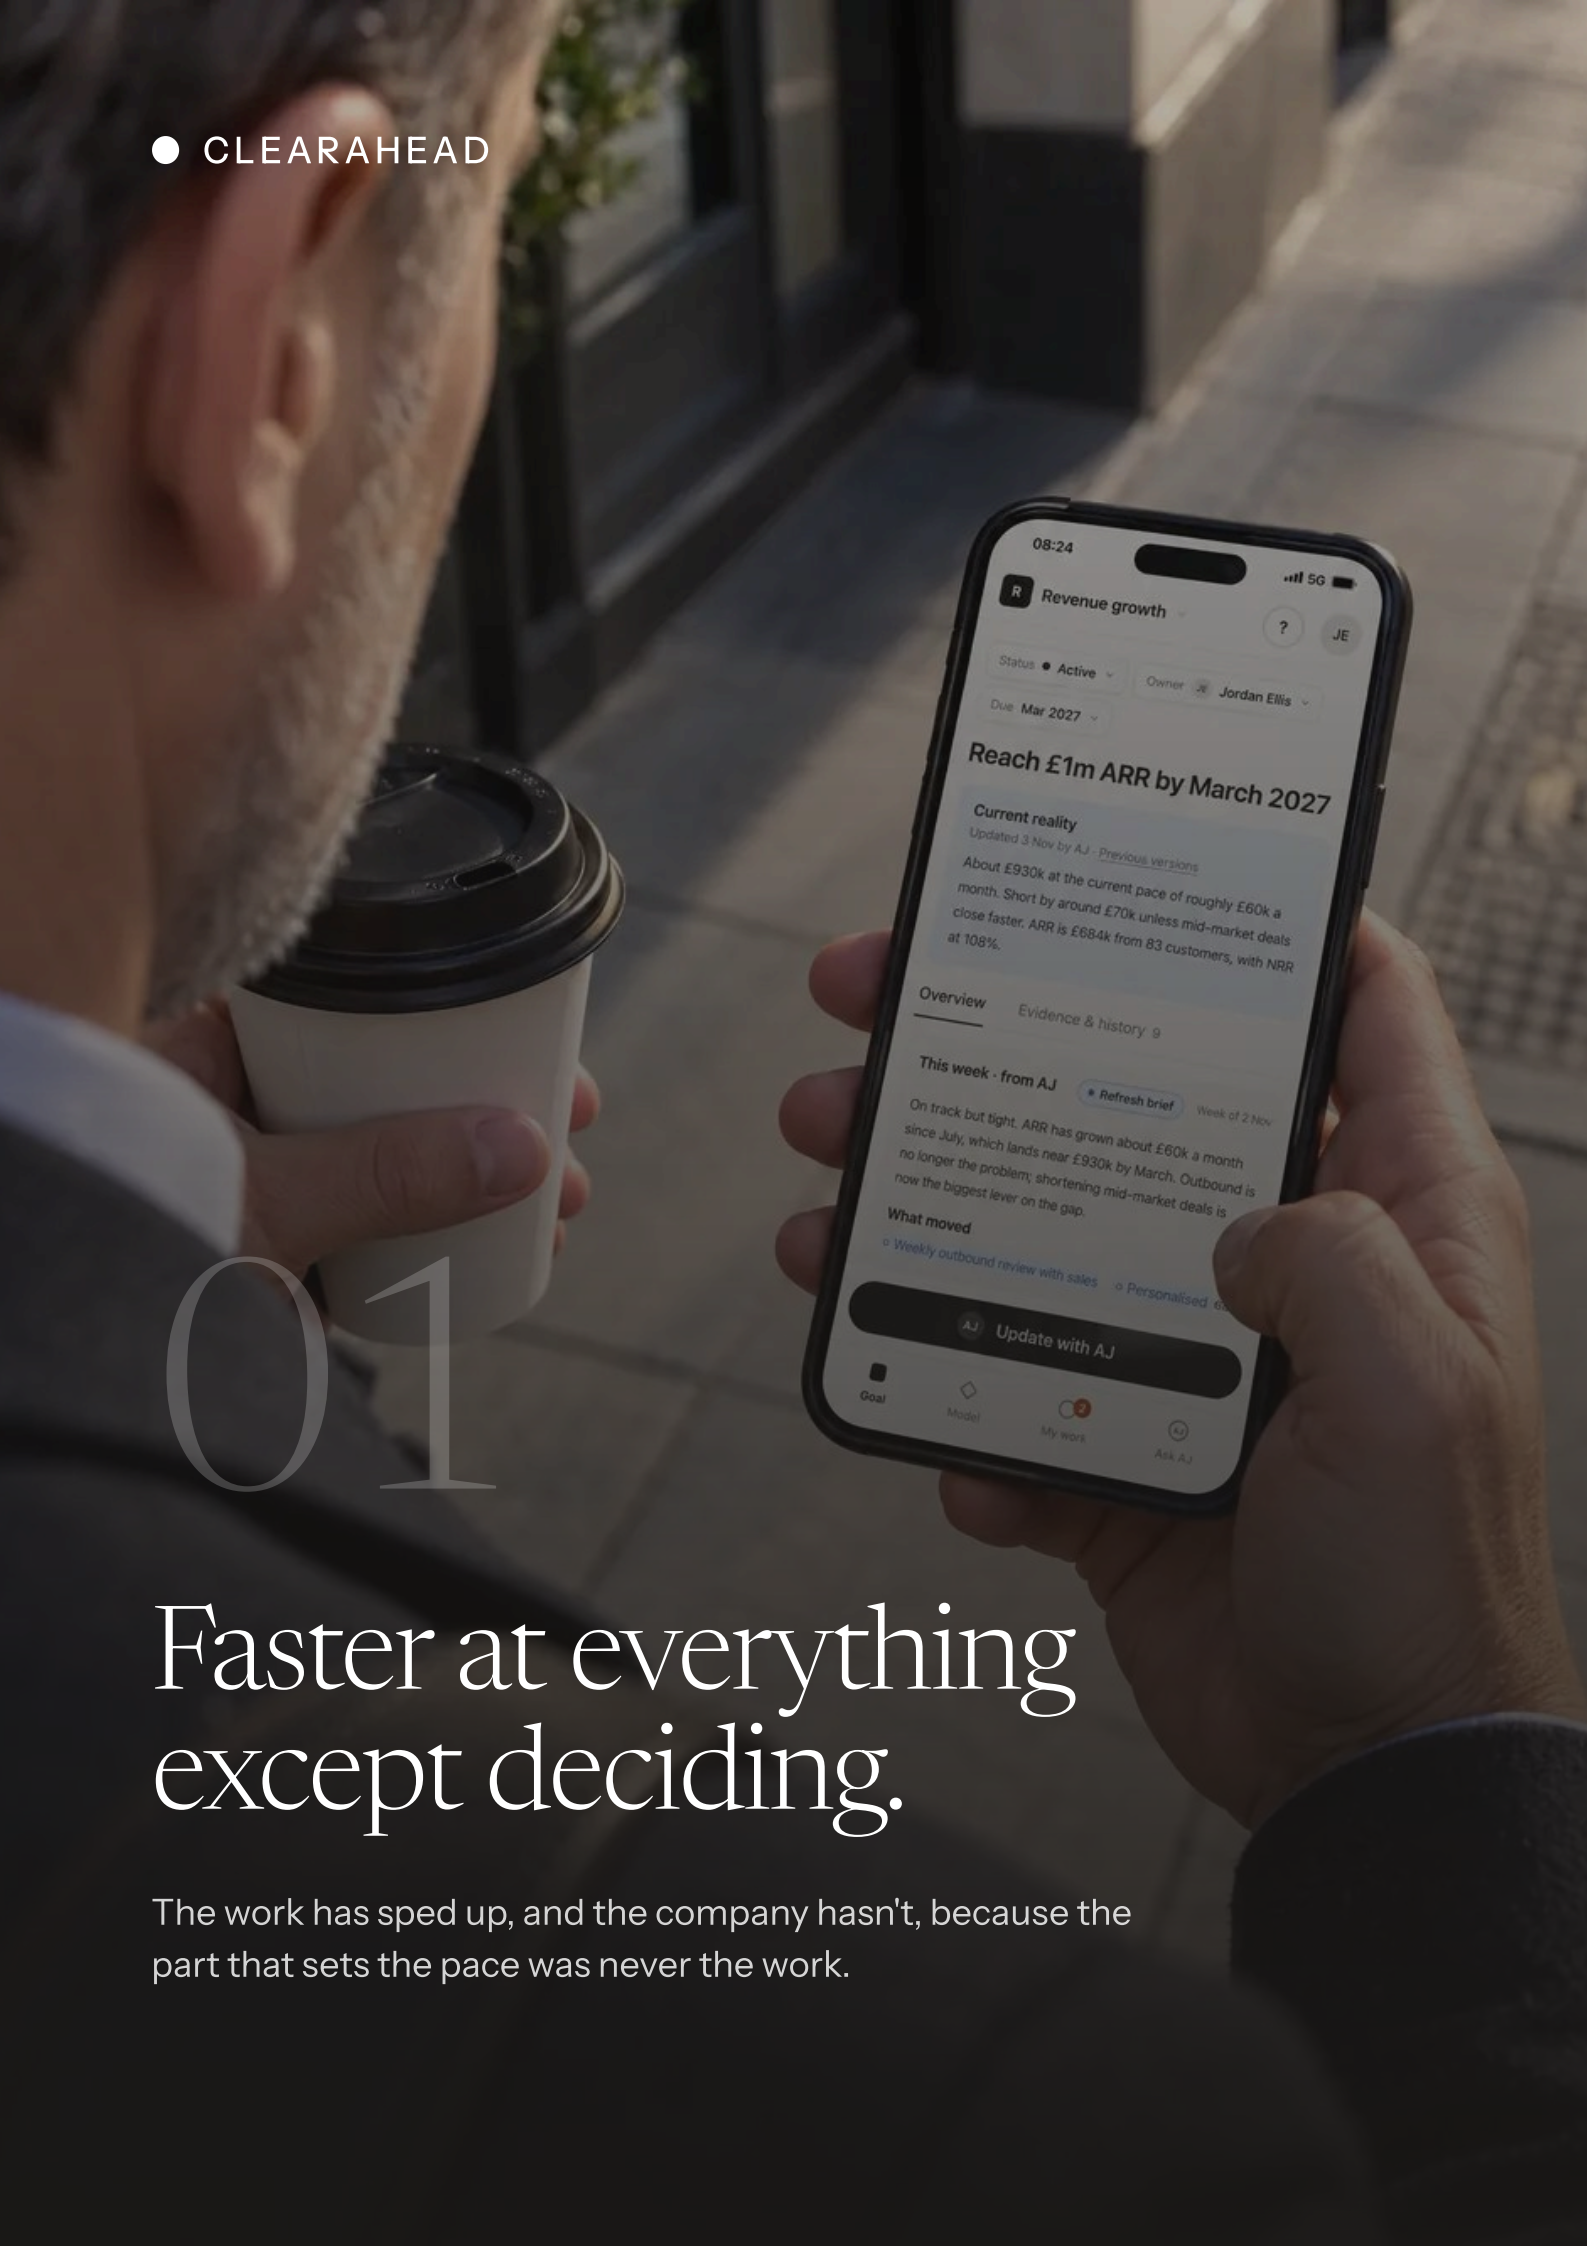
07 A new cadence: the destination is fixed, the route unfolds

08 Goals, constraints, moves and progress

09 Your first thirty days

Sources

CLEARAHEAD



01

Faster at everything except deciding.

The work has sped up, and the company hasn't, because the part that sets the pace was never the work.

Something strange has happened inside the companies that have taken AI seriously. The work itself has sped up, often dramatically, so that a team which used to need a quarter to ship a new pricing page, a sales sequence or the first version of a feature can now do it in a week or two, and yet the company as a whole doesn't feel any faster.

The board pack still arrives once a month, the strategy is still set once a year at an away day, and the big decisions still queue up behind the next planning cycle as though nothing had changed. Leaders feel it as a strange kind of friction; everyone is busy, every dashboard is green, a dozen AI tools have been rolled out, and still the goal they set in January doesn't seem any closer in June.

I think this is the single most important thing that leaders of AI-enabled companies need to understand over the next few years, and it's the argument of this paper: AI has made your company faster at everything except deciding what to do next, which means that for a growing number of businesses the plan itself has become the bottleneck.

I don't mean that as a metaphor. Eliyahu Goldratt, the physicist turned management thinker who gave us the Theory of Constraints, spent his career showing that every system is held back by one constraint at a time, and that improving anything other than that constraint does nothing for the output of the whole.[1] If execution was once the slow part of a company, and AI has made execution fast, the constraint hasn't disappeared; it has moved, and in most companies it has moved upstream, into the heavy, once-a-year machinery of long-term strategy that decides what the fast part should be working on.

The rest of this paper sets out the evidence for that claim, drawing on twenty years of research into strategic planning and the newest studies of what AI does to the speed of work. It then describes the alternative we've built Clearahead around, a cadence where the long-term goal stays fixed as the target, the strategy unfolds as you go, and the whole team works on whatever is holding the goal back right now.

It's a cadence that will feel uncomfortable to some leaders, because there's no eighteen-month Gantt chart at the centre of it, and I understand why that's unsettling. I'll argue, though, that the Gantt chart was failing most businesses long before AI arrived, that the research says so plainly, and that holding on to it now is a choice to let the slowest part of the company set the pace for everyone else.

02

A plan that was out of date before it started

> The company in this section is a composite. It's built from the patterns described in the studies cited in this paper, rather than from any one business, and the details are illustrative.

A composite: Halden, a 60-person B2B software company

In October, Halden's leadership team spends two days at a country hotel building the 2027 plan. They agree to grow annual recurring revenue from £4m to £7m, set five strategic pillars beneath that number, and fill each pillar with initiatives, twenty-three in all, each with an owner, a budget and a set of quarterly milestones. By the end of the second day there's a roadmap on the wall that runs to December 2027, and everyone leaves feeling clear and a little proud, because it's a good plan built by capable people.

The plan is presented to the whole company in January, and within a month it has started to come apart in ways nobody in the room could have seen. The support team rolls out an AI agent that resolves a large share of first-line tickets, so the two hires the plan had pencilled in for the second quarter no longer make sense. Marketing uses AI to produce, in a fortnight, the content programme the plan had scheduled across the whole of the spring. A competitor launches an AI-native product at half Halden's price, and the average sales cycle quietly stretches by three weeks.

None of this is a disaster, and most of it is good news, but none of it is in the plan. By April, the roadmap is something people refer to with a slightly embarrassed smile, most initiatives are marked green because the work on them is getting done, and revenue is still behind. When the leadership team finally digs in, they find that deals are stalling at procurement, because enterprise buyers now ask detailed questions about how Halden uses AI with customer data, and nobody owns the answer. That single issue is holding back more revenue than the other twenty-three initiatives are adding, and it isn't on the chart anywhere.

The decision to fix it waits for the quarterly review in July, because that's when the plan gets revisited, and so the company loses another ten weeks on the one thing that matters most, while every team works hard and on time on things that matter much less.

None of the people at Halden did anything wrong, and that's the point. They planned the way they'd been taught to plan, which is to predict the year, break it into initiatives and then track whether the initiatives are delivered. The trouble is that the method assumes the world will hold still long enough for the prediction to stay useful, and in an AI-enabled company it doesn't; the cost of doing things, the time it takes and even which things are worth doing at all can change inside a single quarter.

What Halden actually needed was much simpler. It needed a clear goal, which it had, and then a way for the whole team to see, week by week, what was really holding that goal back and to move on it straight away, without waiting for the calendar to give them permission. The plan got in the way of that, because it pointed everyone's attention at the initiatives that had been written down rather than at the constraint that had appeared since.

I wrote about the ritual Halden went through, the autumn away day and the plan that's out of date by February, in A loop, not a plan. The research suggests the pattern is far older and far more common than most leaders realise.

03

Long-range planning was failing before AI arrived

It's tempting to blame AI for breaking strategic planning, but the evidence says the patient was already unwell; AI has simply removed the slack that used to hide the symptoms.

In 2005, Michael Mankins and Richard Steele published a study in Harvard Business Review of how well large companies turned their strategies into results. Across the companies they studied, strategies delivered only 63% of the financial value they had promised, and the loss came from the gaps between planning and doing, from resources that never arrived, from priorities that were never made clear and from performance that was never properly tracked against the plan.[2]

A year later, the same authors went further. Working with the Economist Intelligence Unit, they surveyed 156 senior executives at companies with sales of a billion dollars or more, and what they found was a planning process that most leaders had quietly stopped believing in. Only 11% of those executives strongly believed that strategic planning was worth the effort, two-thirds described it as a periodic event, and yet every one of them said they made strategic decisions continuously, without regard to the planning calendar.[3]

63%

of a strategy's promised financial value is what most companies actually deliver.

Mankins & Steele, HBR, 2005

23%

of executives say major strategic decisions are made inside the formal planning process.

McKinsey survey of 796 executives, 2006

28%

of the managers executing a strategy can list three of its priorities.

Sull et al., MIT Sloan Management Review, 124 organisations

The most striking number in that study is the one about decisions. Companies that followed the traditional annual, business-unit planning process made an average of 2.5 major strategic decisions a year, where a major decision was one with the potential to raise profits by 10% or more. Companies that planned continuously, picking up issues as they arose and deciding on them as they went, made 6.1, which is more than twice as many.[3] The annual plan wasn't just slow; it was actively limiting how often the company changed course.

McKinsey found the same thing from a different angle. In a 2006 survey of 796 executives at organisations with revenues of at least $500 million, only 45% said they were satisfied with their strategic planning process, and only 23% said that major strategic decisions were actually made within it. Nearly half said the process failed to track the execution of the strategic initiatives it had produced.[4] In other words, the plan was neither where the decisions happened nor where anyone checked whether they were working, which leaves you wondering what it was for.

Then there's the question of whether anyone knows what the strategy is. Donald Sull and his colleagues analysed 124 organisations for MIT Sloan Management Review and found that only 28% of the executives and middle managers responsible for executing strategy could list three of their company's strategic priorities. Agreement on those priorities fell from 51% among the top team to 22% among the people who reported directly to them, and top teams consistently rated their own company's alignment far higher than everyone else did.[5]

The plan was neither where the decisions happened nor where anyone checked whether they were working, which leaves you wondering what it was for.

None of this is new thinking. Henry Mintzberg argued in 1994, in The Rise and Fall of Strategic Planning, that formal planning can't create strategy at all; it can only programme a strategy that already exists, and real strategy tends to emerge from what a company learns as it acts.[6] Meanwhile the environment those plans are written for has become far less forgiving. Research from the BCG Henderson Institute, published in Harvard Business Review in 2016, found that US public companies were dying at faster and faster rates and had a one in three chance of being delisted within five years.[7]

So the picture before AI was already clear enough: the long-range plan was built once a year, believed in by a small minority of the people who built it, ignored when the real decisions were made, understood by fewer than a third of the people expected to deliver it, and rarely checked against what actually happened. Companies kept doing it because the alternative felt riskier, and because, in a world where execution was slow, a slow plan didn't cost much. The work couldn't move much faster than the plan anyway, so the plan's slowness was hidden inside the general slowness of everything else.

That's the slack AI has removed, and it's why a process that was merely inefficient a decade ago has become, for AI-enabled companies, the thing that sets the speed limit for the whole business.

CLEARAHEAD

04

What AI changed.

The gap between having an idea and having something real has collapsed, and the plan was never designed for that.

The speed of AI isn't hype any more; it's one of the best-measured trends in technology, and the measurements all point the same way.

METR, a research organisation that evaluates AI systems, has tracked the length of task that frontier AI agents can complete on their own. Over roughly six years, that length has doubled about every seven months, and on some software benchmarks the doubling has been faster still.[8] A doubling every seven months means the capability available to your team will roughly double, and then nearly double again, inside a single two-year strategy cycle, which is something no planning process I've ever seen accounts for.

Stanford's AI Index tells the same story from several directions. On SWE-bench, a test of whether AI can solve real software engineering problems, AI systems solved 4.4% of problems in 2023 and 71.7% in 2024. The cost of running a model at the level of the original ChatGPT fell more than 280-fold between late 2022 and late 2024, and the share of organisations reporting that they use AI jumped from 55% to 78% in a single year.[9]

4.4 → 71.7%

of real coding problems solved by AI, from 2023 to 2024.

Stanford AI Index, 2025

280×

fall in the cost of running a GPT-3.5-level model in two years.

Stanford AI Index, 2025

78%

of organisations used AI in 2024, up from 55% the year before.

Stanford AI Index, 2025

Inside companies, the effect on individual work is now well documented. In a pre-registered experiment run by Harvard Business School researchers with Boston Consulting Group, 758 consultants given access to AI completed 12.2% more tasks, finished them 25.1% more quickly and produced work of significantly higher quality, as long as the tasks sat inside what the authors called AI's "jagged frontier". On a task deliberately chosen to sit outside that frontier, consultants using AI were 19% less likely to reach the right answer.[10] Erik Brynjolfsson, Danielle Li and Lindsey Raymond studied more than five thousand customer support agents and found that an AI assistant raised the number of issues resolved per hour by 14% on average, and by 34% for newer and less experienced staff.[11]

Put those findings side by side and the implication for strategy is hard to avoid. Any plan that estimates how long things will take, what they'll cost and how many people they'll need is built on assumptions that are now going out of date in months, so a twelve-month roadmap written in October is describing a company with last autumn's tools, last autumn's costs and last autumn's

idea of what's possible. The initiatives that were scheduled for the third quarter might be done in the first, a hire that made sense in the budget might be unnecessary by spring, and an option nobody considered at the away day might have become cheap enough to be obvious.

The jagged frontier matters too, because it means AI's impact is uneven and hard to predict in advance. Nobody can sit in a room in October and work out which parts of the business AI will transform by June and which it will leave alone, which is why the only sensible approach is to find out as you go and adjust quickly when you do.

A twelve-month roadmap written in October is describing a company with last autumn's tools, last autumn's costs and last autumn's idea of what's possible.

There's a second, subtler effect. When the work was slow, the gap between deciding something and seeing the result was long enough that a quarterly review could keep up, because there wasn't much new information to review. When the work is fast, the evidence arrives continuously, every week brings new results, new surprises and new options, and a company that only looks up once a quarter is ignoring most of what it learns. I wrote in Information is cheap now. Decisions aren't that AI has made gathering and moving information nearly free; what it hasn't done, and can't do on its own, is turn that information into decisions that a whole team understands and acts on.

05

The bottleneck moved

If AI is this powerful, the obvious question is why so many companies that have adopted it aren't seeing the results, and the answer is one Goldratt would have recognised immediately.

In the summer of 2025, MIT's Project NANDA published a report called The GenAI Divide, based on a review of more than three hundred public AI initiatives, 52 interviews and 153 survey responses from senior leaders. Its headline finding was that around 95% of organisations were getting no measurable return on their investment in generative AI, and that only about 5% of integrated AI pilots were extracting real value.[12] The researchers were clear that the problem wasn't the quality of the models. They pointed instead to what they called a learning gap, to tools that don't retain feedback or adapt to context, and to AI that was poorly fitted to the way the company's work actually ran from day to day.

That finding is often read as a story about AI failing, but I read it as a story about constraints. Goldratt's best-known book, The Goal, makes the point with a factory: an hour lost at the bottleneck is an hour lost for the whole system, while an hour saved anywhere else is a mirage.[1] Most companies have spent the last two years saving hours everywhere except the bottleneck. They've made writing faster, coding faster, analysis faster and research faster, and they've left the way they decide, align and learn exactly as it was, running on an annual plan, a monthly board pack and a weekly status meeting.

An hour lost at the bottleneck is an hour lost for the whole system, while an hour saved anywhere else is a mirage.

Eliyahu M. Goldratt, The Goal

Goldratt made exactly this argument about the last great wave of business technology. In Necessary But Not Sufficient, written about the enterprise software boom of the 1990s, he asked what limitation a new technology actually removes and whether the company had changed the rules it used to live with that limitation.[13] If it hadn't, the technology was necessary but not sufficient, and the benefits never arrived. I applied his four questions to AI in AI is necessary but not sufficient, and the conclusion there is the same as here: AI removes the limitation of slow, expensive execution, but most companies are still living by the rules that slow execution created, including the rule that strategy is something you set once a year and then deliver.

So where is the bottleneck now? In most AI-enabled companies I talk to, it sits in three places that are really one. The first is deciding, because the queue of decisions waiting for the next planning cycle, the next board meeting or the next leadership offsite is longer than ever. The second is aligning, because when every team can move fast, the cost of teams moving fast in different directions goes up sharply. The third is learning, because a company that only reviews its strategy once or twice a year can't absorb what its own fast work is teaching it, which is exactly the learning gap the MIT researchers described.

The good news is that the Theory of Constraints has a strong track record when it's applied seriously. Vicky Mabin and Steven Balderstone reviewed more than eighty published applications of Goldratt's methods and found significant improvements in both operational and financial performance.[14] The method works because it refuses to spread effort across everything; it asks what's holding the system back right now, puts the effort there, and then asks the question again. That's precisely the discipline an AI-enabled company needs at the level of strategy, and it's the opposite of what a long-range plan with twenty-three initiatives encourages.

I've written more about how this thinking shaped Clearahead in Remove constraints, don't add initiatives and The five focusing steps, run on a whole company.

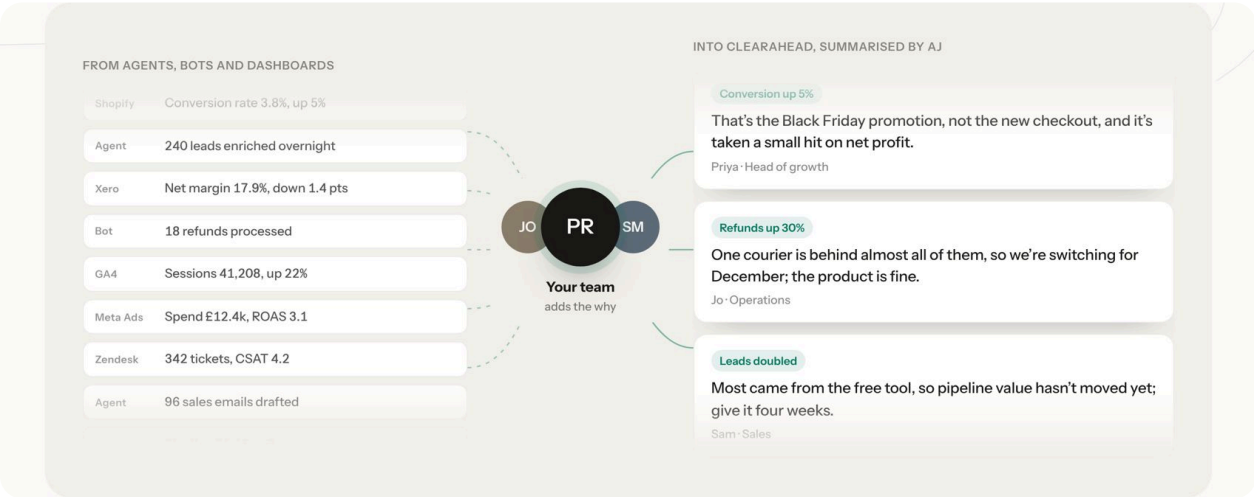
06

Why leaders hold on to the Gantt chart

If the evidence is this clear, it's worth asking why the long-range plan has survived, and the honest answer is that it does a real job for leaders, just not the job it claims to do.

A long-term plan with a Gantt chart at the centre gives a leadership team something it badly wants, which is a feeling of control. It turns an uncertain future into a set of bars on a timeline, each with an owner and a date, and it gives the board, the investors and the team a story they can understand and a way of holding people to account. When a leader suggests dropping it, the fear isn't irrational; it's the fear of standing in front of a board or a team without a map and admitting that you don't know exactly how the next eighteen months will go.

I take that fear seriously, because I've felt it, and because the people who feel it are usually the ones who care most about delivering. But the research in this paper makes a hard point, which is that the control was mostly an illusion even before AI. If strategies deliver 63% of their promised value, if fewer than a quarter of major decisions are made inside the process, and if fewer than a third of the people delivering a plan can say what it is, then the chart was offering certainty it couldn't back up. The future was always uncertain; the plan just made that uncertainty easier to ignore until the quarterly review.



The goal stays fixed; the route is visible to everyone. Clearahead shows one goal, what's in its way and who is moving on each part of it, updated as the work happens.

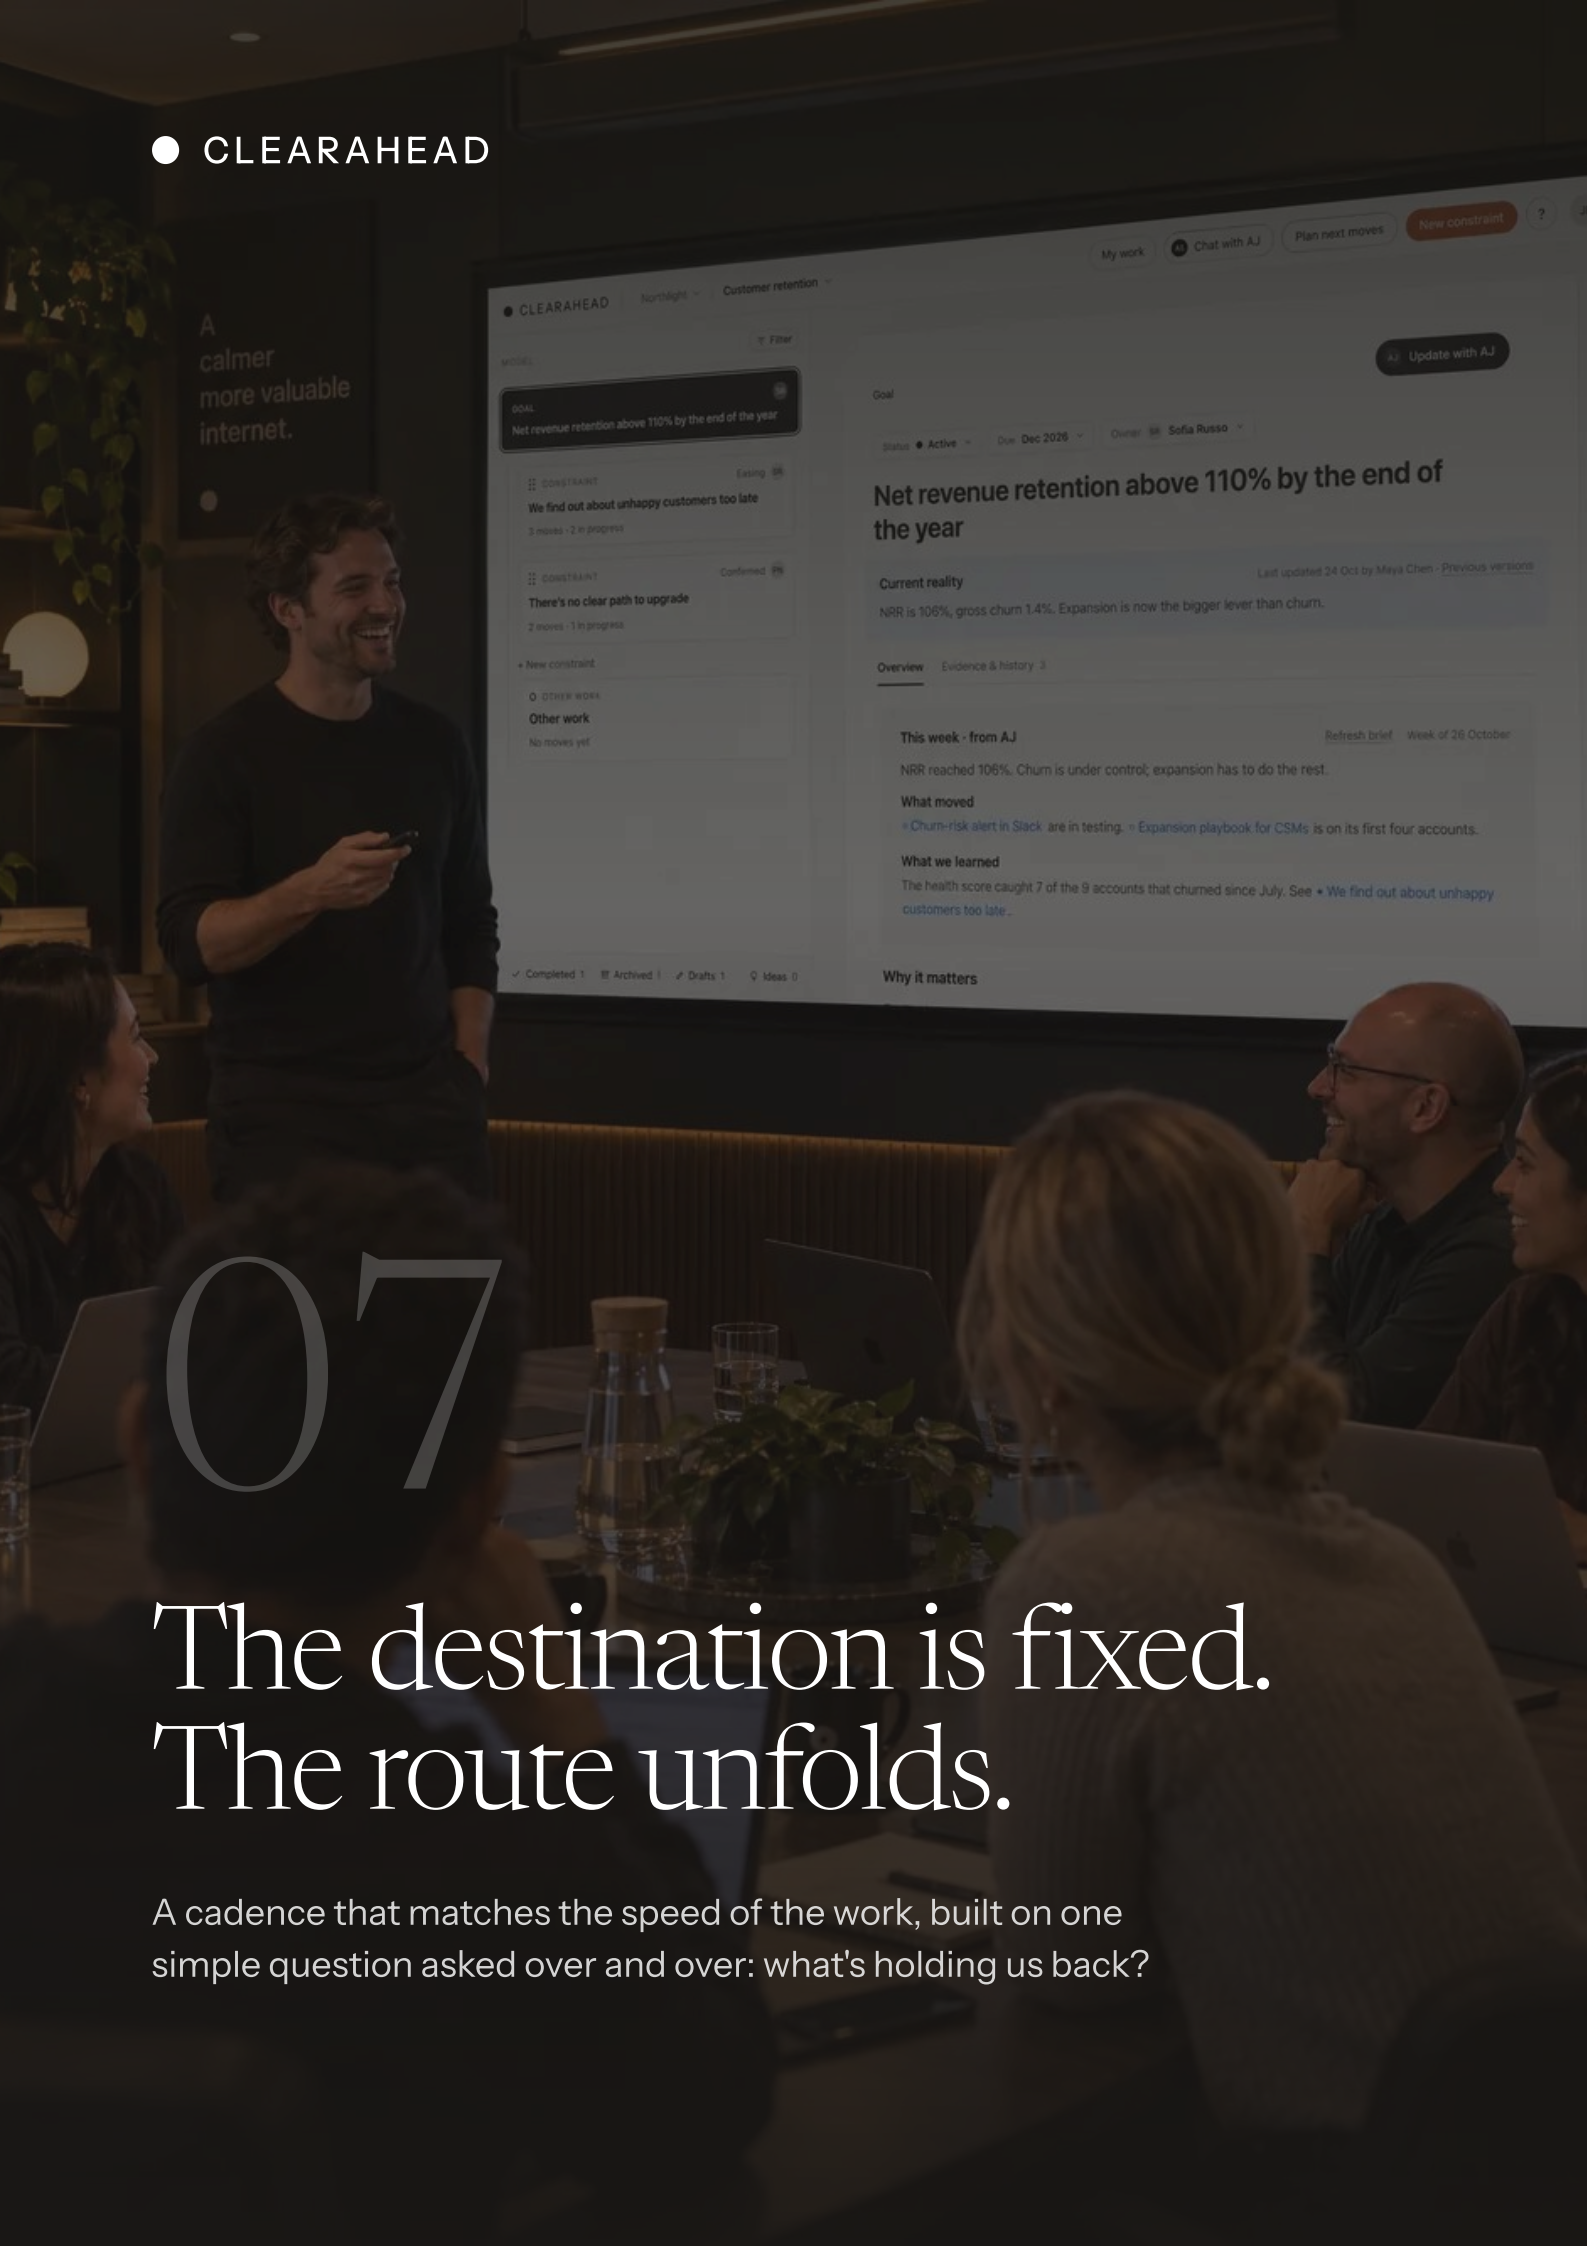
It helps to separate what leaders actually need from the artefact they've used to get it. They need direction, so that everyone knows where the company is going; accountability, so that everyone knows who owns what; visibility, so that they can see whether things are working; and confidence, so that they can tell a board or a team that the company is on track and mean it. A Gantt chart offers all four on the day it's drawn, and then each one decays as the chart drifts away from what's really happening.

Every one of those needs can be met without predicting the next eighteen months, and met better. Direction comes from a clear long-term goal that doesn't change every quarter. Accountability comes from every piece of work having one owner and a due date, which works just as well when the work is planned weeks ahead as when it's planned a year ahead. Visibility comes from a live picture that updates as the work moves, rather than a status report assembled once a month, and confidence comes from seeing, every week, that the thing most in the way of the goal is being dealt with.

There's a line usually credited to Dwight Eisenhower that plans are worthless but planning is everything, and I think it gets to the heart of this. The thinking that goes into a strategy, the hard conversations about where you're going and what's in the way, is enormously valuable, and nobody should stop doing it. What's lost its value, at least in AI-enabled companies, is the document that comes out of that thinking and the rituals that grow up around defending it.

As for boards and investors, in my experience they would far rather see a live picture of the goal, what's holding it back and what's being done about it, updated every week, than a beautifully formatted plan that everyone in the room privately knows is already out of date. A board that can read along with the company's progress asks better questions, and a leadership team that isn't defending a stale plan gives better answers.

CLEARAHEAD



07

The destination is fixed. The route unfolds.

A cadence that matches the speed of the work, built on one simple question asked over and over: what's holding us back?

If the plan is the bottleneck, the answer isn't to plan faster, and it certainly isn't to stop thinking about the long term; it's to change what you fix in advance and what you let unfold.

The model we've built Clearahead around keeps one thing fixed and lets everything else move. The thing that's fixed is the goal, the long-term outcome the company is working towards, written in the team's own words with an owner and a date, whether that's reaching £10m in annual recurring revenue by the end of 2028, becoming the default choice for a particular kind of customer, or getting a product to profitability. The goal is the target, and it should be stable enough that everyone in the company can remember it and steer by it.

What moves is the route. Instead of deciding in advance every initiative that will get the company to its goal, the team asks one question, again and again: what's holding us back from this goal right now? The answers are the constraints, a short list of the real things in the way, put in order by the team so that the one at the top gets the most attention, and under each constraint sit a few moves, the specific actions that will clear it, each with an owner and a due date. When a constraint clears, the team records what cleared it, the next one moves up, and the question gets asked again.

The strategy is the sequence of constraints you clear on the way to the goal, and you only find out what that sequence is by clearing them.

In this model, the strategy isn't a document you write at the start; it's the sequence of constraints you clear on the way to the goal, and you only find out what that sequence is by clearing them. That sounds radical, but it's how most good strategy has always been made in practice, which is what Mintzberg meant by emergent strategy, and it's what the executives in the Mankins and Steele study were describing when they said they made strategic decisions continuously, issue by issue, whatever the planning calendar said.[3, 6] The difference is that it now happens in the open, with the whole team able to see it, rather than informally and out of sight.

This doesn't throw away the long-term thinking a company already does. The vision, the three-year picture and the annual priorities still matter, because a team needs to know where it's going before it can tell what's in the way. What changes is that the long-term picture sets the destination and the loop keeps the plan alive, so that instead of a plan written once and revisited at the next away day, the team works on whatever is holding it back today, learns as it goes and adjusts the route as the market, the people and the numbers change.

The rhythm

A loop needs a rhythm, and the one we suggest is designed around what people actually do in Clearahead and what AJ, the AI built into it, does on a schedule. People only meet when a decision needs the team together, and teams fit the rhythm into the cadence they already have rather than adopting a new set of meetings.

As it happens	People tell AJ what's changed in a sentence when the work moves, and AJ updates the current reality of each goal, constraint and move, so the picture is always current.
Every week	AJ writes a brief for the team first thing on Monday, covering where the goal stands, what moved, what didn't and which decisions need someone, and the team takes the decisions that need taking.
Every month	The team looks at the constraints as a whole, asks whether the order is still right and whether anything new has appeared, and reorders them if it has.
Every quarter	The team steps back and asks whether the goal itself is still the right one, and if it has changed, changes it and re-checks every constraint against the new version.
Turning points	When a constraint clears or the goal is reached, AJ writes up what happened and what cleared it, so the company builds a memory of what works.

Compare that with the annual cycle in section three. The decisions that used to wait for the planning calendar now get taken in the week they're needed, which is what the continuous planners in the Mankins and Steele study were doing when they made more than twice as many major decisions a year. The learning that used to be lost between reviews now gets captured every time a constraint clears, which is a direct answer to the learning gap the MIT researchers found. And because everyone works from the same live picture of the goal and what's in its way, alignment stops being something a leadership team hopes it achieved at the away day and becomes something the whole company can see every week.

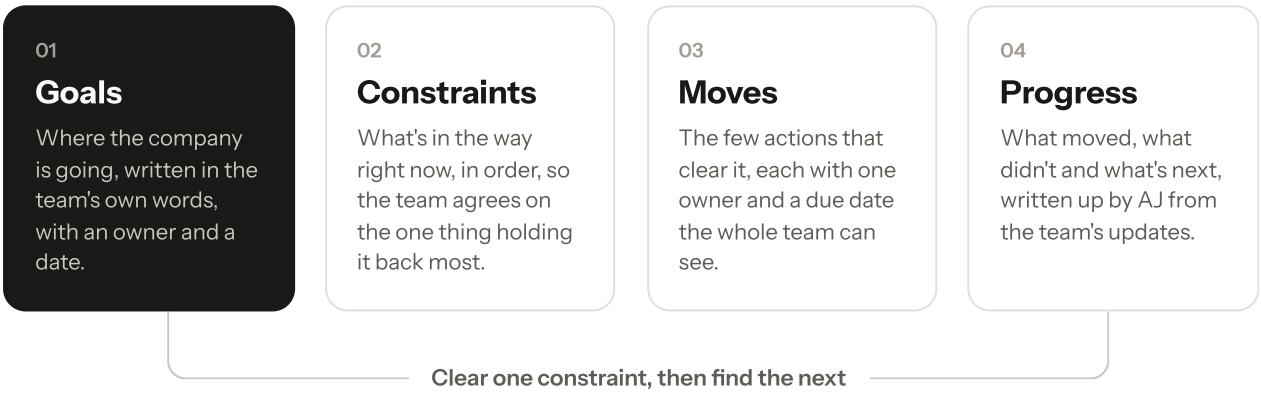
I set out the whole operating model, including the principles behind it, in The operating model for the AI era, and the reasons alignment matters more than effort in Alignment beats activity.

08

Goals, constraints, moves and progress

Clearahead is goal tracking for leadership teams, built on four parts and one loop, with AJ, our AI, keeping the whole thing current so that the team can spend its time on the thinking and the deciding.

I want to be clear from the start that AJ is an AI, because it does a lot of the work that used to need meetings and reports, and because the way it does that work matters. AJ asks before it answers, challenges rather than prescribes, and leaves every decision with the people who own it; I've written about why in The machine asks, people decide and Why AJ asks instead of answers.



The long-term plan sets the destination.

The vision, the three-year picture and the annual priorities still matter, because a team needs to know where it's going before it can tell what's in the way.

The loop keeps the plan alive.

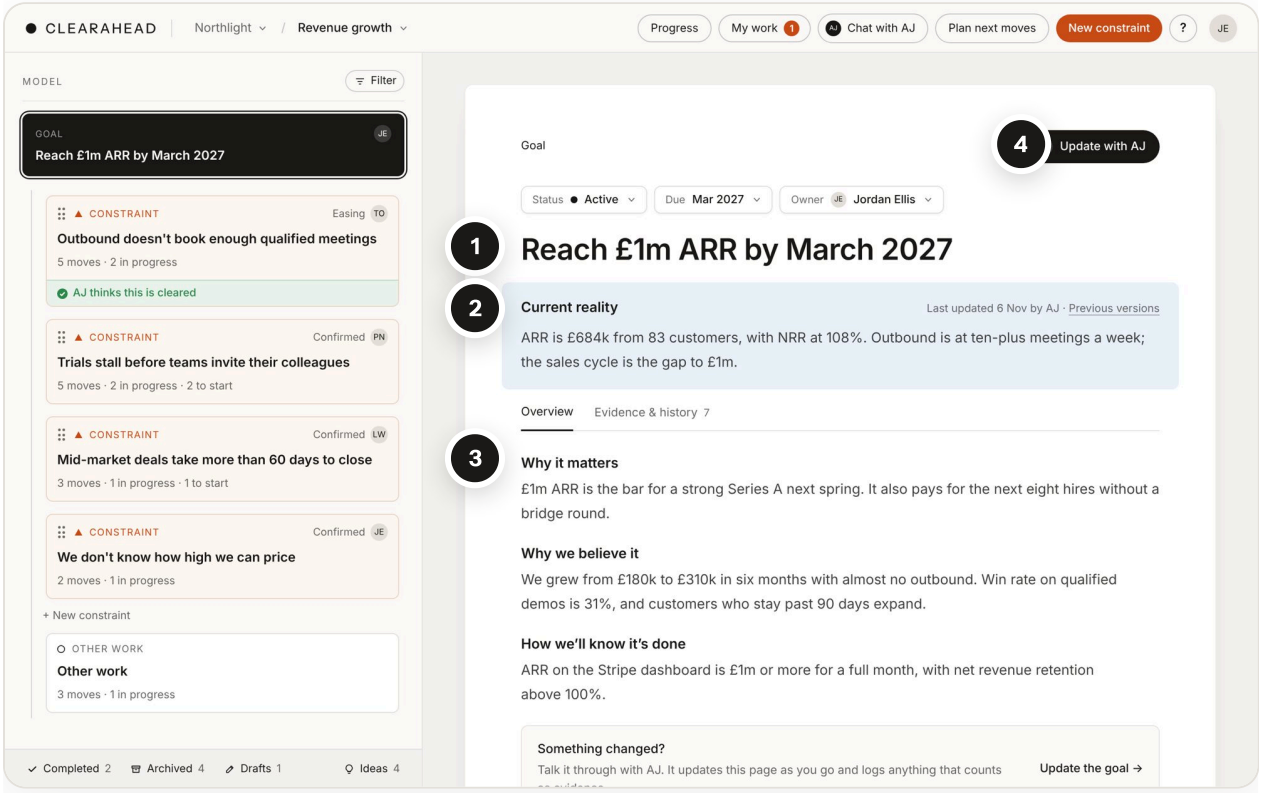
Instead of a plan written once and revisited at the next away day, the team works on whatever is holding it back today, learns as it goes and adjusts the route as things change.

Over the next few pages you can see how each of the four parts works in the app, and how it answers the problems this paper has described.

08 · 01

Goals: the destination stays fixed

Every workspace has one goal at the top, with an owner and a due date, so everything underneath it points the same way. This is the fixed destination the new cadence depends on, written once in the team's own words and kept in view, and because an organisation can run up to ten workspaces, the company goal can sit alongside goals for sales, product or operations, each with its own constraints and moves.



1 **The goal sits at the top.**
One goal with an owner and a due date, written in the team's own words.

2 **Current reality, always current.**
A short, plain summary of where things stand, rewritten by AJ as updates come in.

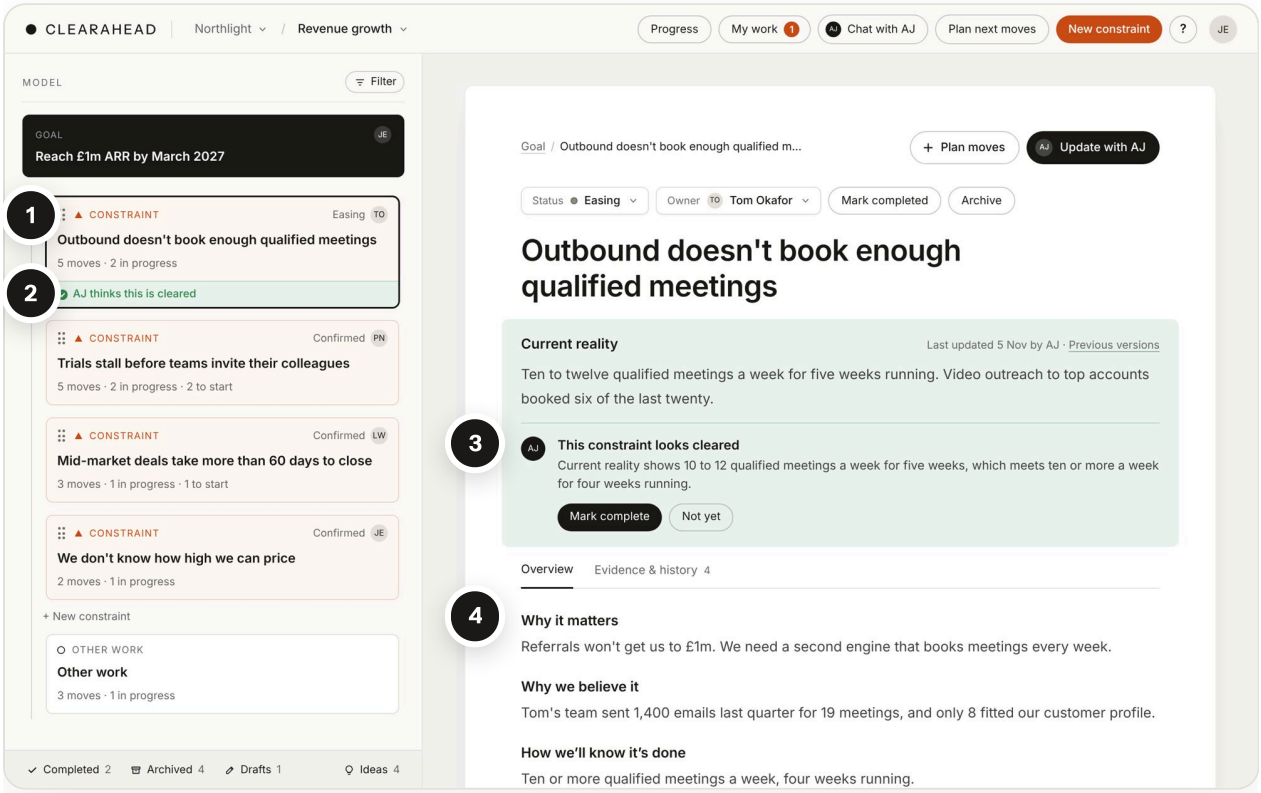
3 **The thinking behind it.**
Why it matters, why we believe it and how we'll know it's done, written once and kept in view.

4 **Update with AJ.**
Tell AJ what's changed in a sentence and the page updates itself while you watch.

08 · 02

Constraints: what's in the way right now

The constraints are what's in the way of the goal right now, and the team puts them in the order that matters, so the one at the top is the one it clears first; we suggest three to five active constraints, which is few enough to hold in your head and enough to see the shape of the problem. When a team sets up a goal, AJ asks the plain question, what's holding you back, and keeps asking until the answers get past the symptoms to the real causes, which is where the Halden team would have found its procurement problem in a week instead of waiting until July.



1 Constraints, in order.
Drag them into the order that matters; the one at the top is what the team clears first.

2 AJ spots when one is cleared.
It checks the team's updates against how you said you'd know it was done.

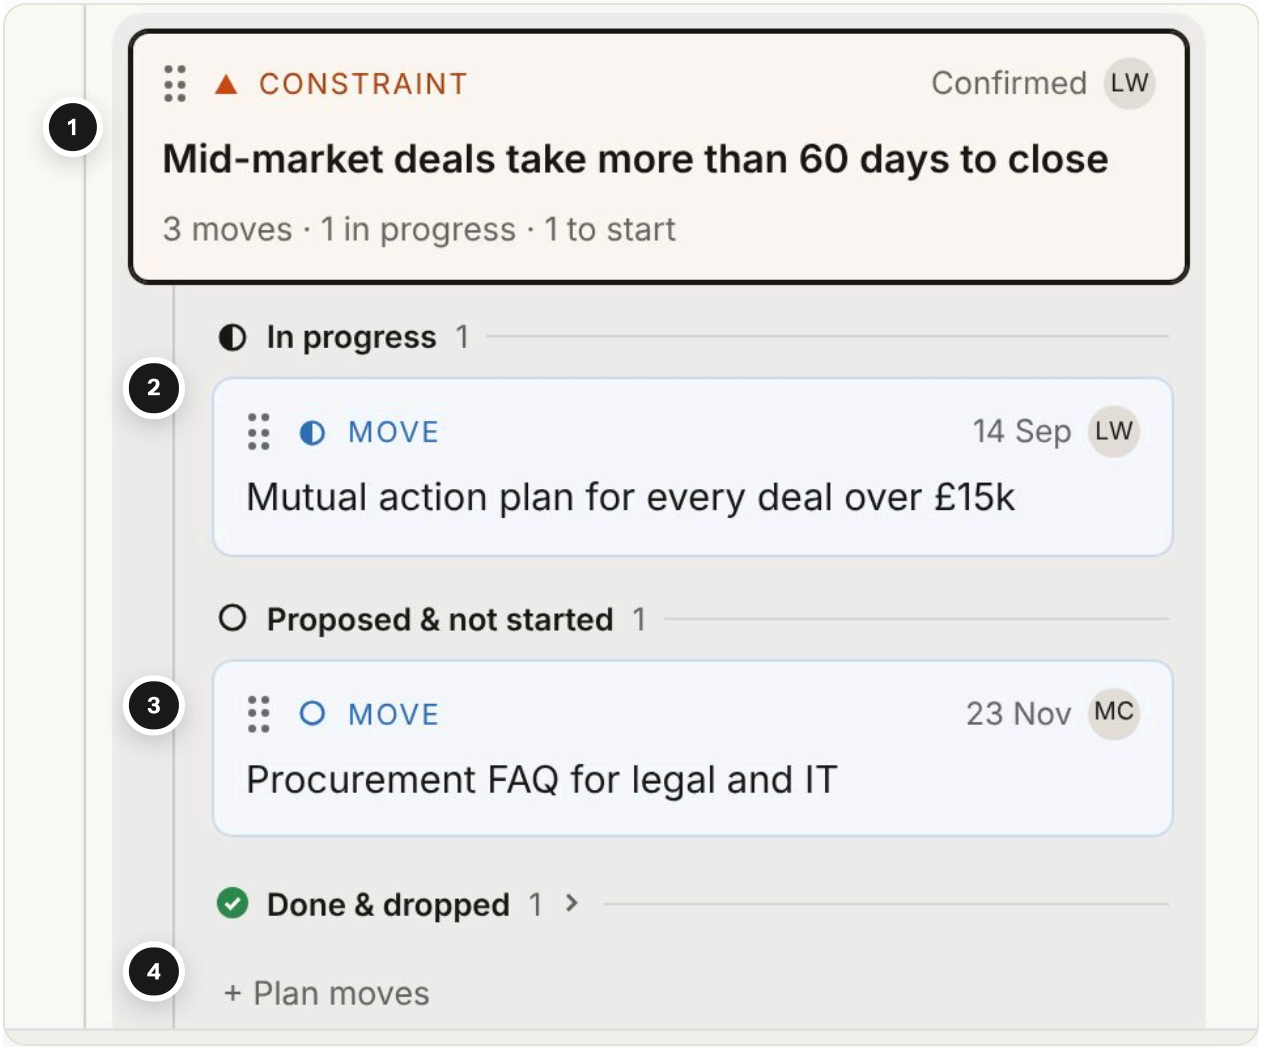
3 AJ suggests, you decide.
The owner marks it complete or says not yet, so the call stays with people.

4 Why it's in the way.
The reasoning behind each constraint stays on the page, so nobody has to ask twice.

08 · 03

Moves: the work that clears it

Under each constraint sit the moves, the few specific actions that will clear it, each with one owner and a due date the whole team can see. AJ suggests moves, the team keeps the ones it believes in and adds its own, and good ideas that don't fit the current constraint go onto a list to come back to. Moves are planned weeks ahead rather than a year ahead, which keeps all the accountability of a roadmap without the prediction, and because every move hangs off a constraint, every piece of work has a reason the whole team can read.



1 Open a constraint.
Its moves open beneath it, so the work always sits under the reason for it.

2 One owner, one date.
Every move has a single owner and a due date the whole team can see.

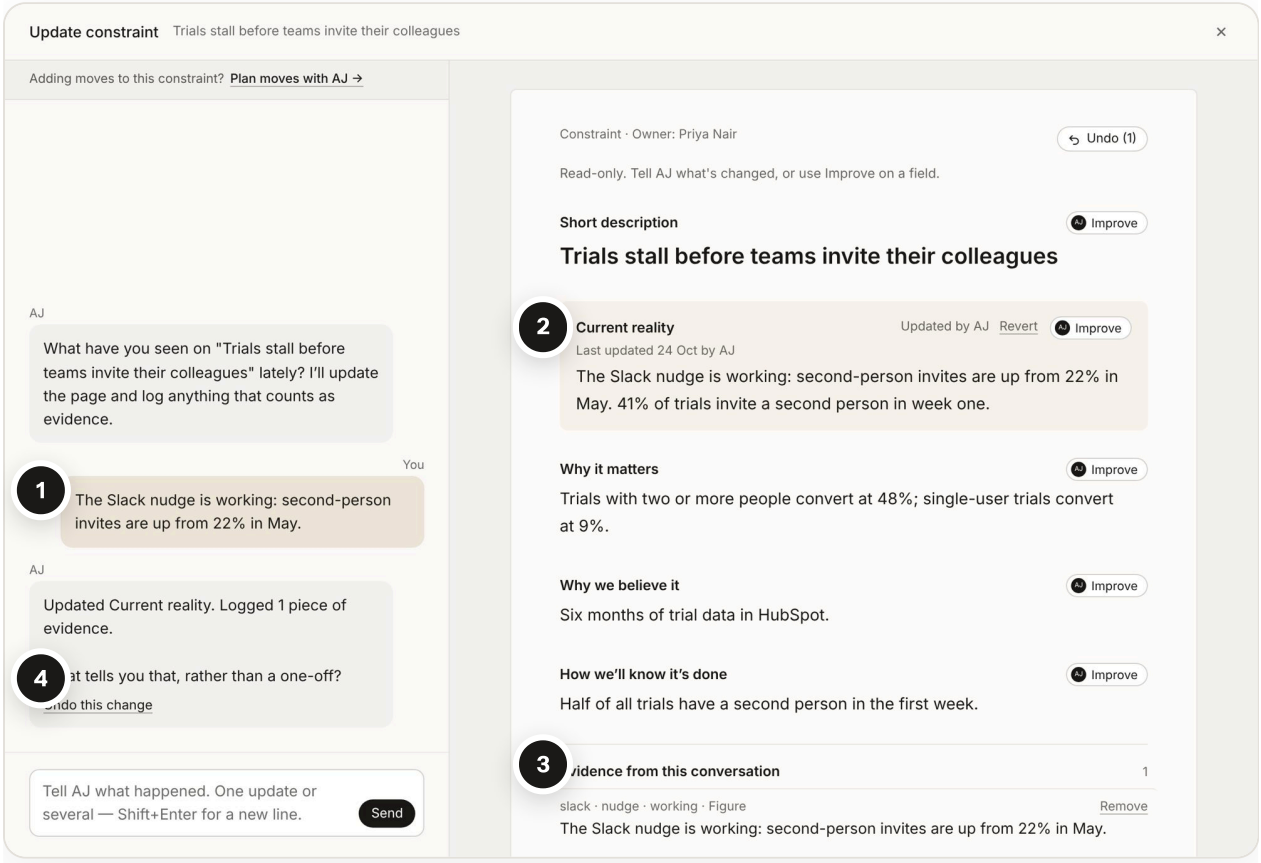
3 Grouped by where they are.
In progress, not started and done, so it's clear at a glance what's moving.

4 Plan moves with AJ.
AJ suggests moves; you keep the ones you believe in and add your own.

08 · 04

Progress: updated with a conversation

People update their work by telling AJ what's changed in a sentence, the way they'd say it in a meeting, and AJ rewrites the current reality, logs anything that counts as evidence and asks the next question, like a good coach would. Every Monday, AJ writes a brief covering where the goal stands, what moved and what needs a decision, and when the goal is reached, AJ writes the story of how you got there, which becomes the company's memory of what worked.



1 Say what happened.
Type an update the way you'd say it in a meeting, one line or several.

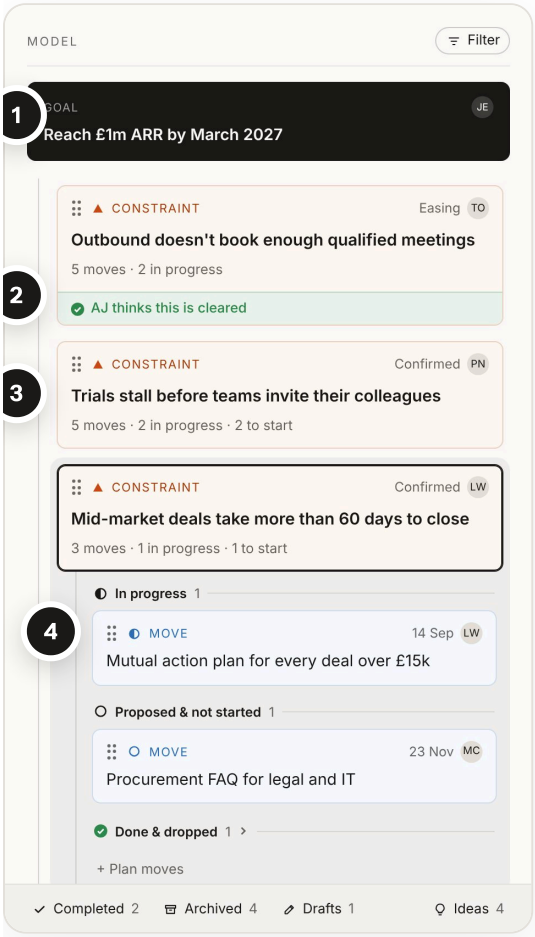
2 AJ updates the page.
It rewrites the current reality straight away, and any change can be undone.

3 Evidence, logged.
Anything that counts as evidence is saved, so the team can see why things moved.

4 Then it asks the next question.
AJ checks whether it's a real shift or a one-off before anyone calls it cleared.

One picture the whole team reads

Put together, the four parts are a working version of the cadence in the last section: the goal gives the company its direction, the constraints keep everyone's attention on the real bottleneck, the moves keep the work accountable without pretending to know the year in advance, and the progress loop turns fast work into fast learning and fast decisions.



1 **The goal sits at the top.**
Every workspace has one goal with an owner and a due date, so everything underneath it points the same way.

2 **AJ spots when a constraint is cleared.**
It checks the team's updates against how you said you'd know, then asks the owner to confirm, so the call stays with people.

3 **Constraints, in order.**
Drag them into the order that matters; the one at the top is what the team clears first, and each card shows its moves and who owns it.

4 **The moves that clear it.**
Open a constraint to see the moves under it, grouped into in progress, not started and done, each with one owner and a due date.

It's built for the whole company rather than for any one person's productivity, because alignment only happens when everyone, from the founder to the board, is looking at the same picture; that's why viewers are free and unlimited, and why the paid plan is one price for the whole company. You can see each part in more detail on How it works.

09

Your first thirty days

One of the quiet advantages of this cadence is that you don't need a six-month implementation to start it, because the model is small enough to set up in an afternoon and the rhythm builds itself as the team uses it.

Week one: set the goal

Spend an hour with AJ setting the goal, with the outcome, a rough date and how you'll know it's done. Let AJ ask what's holding you back until you have three to five real constraints in your order, put a few moves under the top one, and invite the team.

Week two: let the work talk

Ask people to tell AJ what's changed whenever their work moves. Read the first Monday brief together, take the decisions it surfaces, and pick one status meeting or report that the brief has made redundant and drop it for a month.

Week three: take decisions in the week

When something new gets in the way, add it as a constraint and reorder the list there and then, rather than saving it for a review. Treat any decision that's waiting for a calendar date as a warning sign.

Week four: review the order

Look at the constraints as a whole and ask whether the order is still right. If the top one has cleared, record what cleared it and move on to the next, and invite your board or investors to read along.

Clearahead is free for up to three people with no card needed, which suits most early-stage leadership teams, and after that it's one flat price for the whole company however many people you invite. If you'd rather have help, we can walk your team through the first goal, and if you work with a business coach or adviser, they can join your workspace as a free viewer and read along between sessions. I wrote more about why getting started should take an afternoon in Hours, not months.

Stop letting the plan set the pace

The companies that win the next few years won't be the ones with the most AI tools, because everyone will have those; they'll be the ones that change the rules they run by to match what the tools have made possible. For most leadership teams, the most important of those rules is the one that says strategy is something you set once a year and then deliver, because that rule is now the thing standing between a fast company and a fast-moving goal.

So keep the long-term goal, think hard about it and write it down where everyone can see it, and then let go of the idea that you can predict the route. Ask what's holding you back, clear it, learn what cleared it and ask again, every week, with the whole team looking at the same picture. The strategy will unfold as you go, and when you look back, you'll find it was a better strategy than the one you'd have drawn on the wall in October, because it was built on what you learned rather than on what you guessed.

Further reading on clearahead.ai

A loop, not a plan
The operating model for the AI era
Remove constraints, don't add initiatives
AI is necessary but not sufficient
Information is cheap now. Decisions aren't
How Clearahead works



Alex Durston is the founder of Clearahead, goal tracking for leadership teams. He built Clearahead to help companies align around one goal, find what's really in the way and clear it, with AJ, an AI built to ask the questions a good coach would.

alex@clearahead.ai · linkedin.com/in/alexdurston

Sources

The Halden example in section two is a composite, built from the patterns in these studies; it doesn't describe a real company. Statistics are quoted as published by the original authors.

1 Eliyahu M. Goldratt and Jeff Cox, The Goal: A Process of Ongoing Improvement, North River Press, 1984.

2 Michael C. Mankins and Richard Steele, "Turning Great Strategy into Great Performance", Harvard Business Review, July 2005. bain.com/insights/turning-great-strategy-into-great-performance

3 Michael C. Mankins and Richard Steele, "Stop Making Plans; Start Making Decisions", Harvard Business Review, January 2006, reporting a Marakon Associates and Economist Intelligence Unit survey of 156 senior executives. bain.com/insights/stop-making-plans-start-making-decisions

4 Renée Dye and Olivier Sibony, "How to improve strategic planning", The McKinsey Quarterly, 2007 no. 3, reporting a McKinsey survey of 796 executives, 2006.

5 Donald Sull and colleagues, "No One Knows Your Strategy — Not Even Your Top Leaders", MIT Sloan Management Review, 2018. sloanreview.mit.edu

6 Henry Mintzberg, The Rise and Fall of Strategic Planning, Free Press, 1994.

7 Martin Reeves, Simon Levin and Daichi Ueda, "The Biology of Corporate Survival", Harvard Business Review, January–February 2016. hbr.org/2016/01/the-biology-of-corporate-survival

8 Thomas Kwa, Ben West, Joel Becker and colleagues, "Measuring AI Ability to Complete Long Tasks", METR, March 2025. metr.org

9 Stanford Institute for Human-Centered AI, The AI Index 2025 Annual Report, April 2025. hai.stanford.edu/ai-index/2025-ai-index-report

10 Fabrizio Dell'Acqua, Edward McFowland III, Ethan Mollick and colleagues, "Navigating the Jagged Technological Frontier", Harvard Business School Working Paper 24-013, September 2023. hbs.edu

11 Erik Brynjolfsson, Danielle Li and Lindsey Raymond, "Generative AI at Work", Quarterly Journal of Economics 140(2), 2025. nber.org/papers/w31161

12 MIT Project NANDA, The GenAI Divide: State of AI in Business 2025, July 2025.

13 Eliyahu M. Goldratt, Eli Schragenheim and Carol A. Ptak, Necessary But Not Sufficient, North River Press, 2000.

14 Victoria J. Mabin and Steven J. Balderstone, "The performance of the theory of constraints methodology", International Journal of Operations & Production Management 23(6), 2003.

CLEARAHEAD

The modern way to set, track and achieve company goals.

Set the goal, find what's in the way and clear it, with the whole team on one live picture and AJ, our AI, keeping it current.

Start free at clearahead.ai

Free for up to 3 people, no card needed.

Clearahead AI Ltd · Company number 17557849 · 71-75 Shelton Street, Covent Garden, London WC2H 9JQ