

EBOOK

Planning the 2027 workforce

How leading organizations align people, plans, and budgets in a world of constant change.



Align teams

Unify HR, Finance, and Talent.



Plan with confidence

Model scenarios and what-ifs.



Make smarter investments

Allocate resources to drive impact.

TeamOhana is the workforce planning platform that unites Finance, Talent, and HR to plan, approve, and manage headcount and AI spend in one system.

Visit teamohana.com →

Planning the 2027 workforce: people and digital workers on one budget

Every dollar in your 2027 plan can now go to three different things: a new hire, more AI tools for someone already on your team, or a digital worker that does the work instead. Read this before you build the plan, together.

How to use this

Headcount is the one expense line where Finance, Talent, and HR each hold a piece of the decision, and none of you holds the whole thing. This guide works the same way. Read it start to finish, or use the contents below to jump straight to the part your team needs. Every section ends with a short breakdown for Finance, Talent, and HR, so you can read your part fast or read it all for the full picture.

BEFORE YOU PLAN

Share this with your Finance, Talent, and HR partners before you sit down to build the plan. Run the readiness score in Part 11 out loud, together, before you write a single line. It's built as a printable checklist right in this document, check the boxes as a group. If the three of you disagree on the answers, that disagreement is the most useful thing that happens in the room.

Contents

- 01** Your 2027 plan has to price people and AI spend as one budget
- 02** A hire, an AI upgrade, or a digital worker now compete for the same dollar
- 03** Five ways the 2027 plan breaks
- 04** What planning season actually costs
- 05** How do you manage AI token spend alongside your headcount budget?
- 06** TeamOhana Computer: your AI coworker for planning
- 07** The ledger for AI spend
- 08** The planning sequence
- 09** What happens after the plan locks
- 10** Where this sits in your stack

11 The readiness score

12 What changes when this is one system

13 What this looks like in practice

14 Next step

Your 2027 plan has to price people and AI spend as one budget

2027 planning asks you to redesign work, not forecast headcount

Every planning cycle before this one asked the same question: how many people, and at what cost. 2027 asks something harder. Which work needs a person, which work needs a person with better tools, and which work needs neither. Nobody in your company has answered that at scale before, and the answer changes org charts, job descriptions, and career paths. Treat this as the first step of a workforce redesign that runs across the company, because that is what it is.

Boards now expect an approval trail for every headcount decision

That tolerance disappeared with cheap money. Boards want to know which roles were approved, by whom, at what cost, and what happened to the ones that were funded and never filled. Those are controls questions, and most companies cannot answer them from a system.

AI spend is growing too fast to track after the fact. It has to be budgeted upfront.

Enterprise AI spend is growing fast enough that it can no longer sit outside the budget conversation. AI-native application spend rose 108 percent year over year across all company sizes, and 400 percent at organizations above 10,000 employees. Vendors are already converting seat pricing into token and action based pricing. [Gartner recorded the demand side response](#): CFO headcount growth expectations falling from 6 percent to 2 percent between 2025 and 2026, which its analysts describe as a structural pivot from labor expansion to automation, not a side experiment.

Other companies have already named one owner for AI and headcount together

This is not only a TeamOhana argument. [Atlassian recently gave its Chief People Officer responsibility for AI enablement](#), expanding her remit from a 700 person HR organization to a 3,500 person function. Their stated reason: cultural and technological transformation cannot be separated if the change is going to stick. Name one executive who owns the mixed workforce decision at your company, and do it before planning opens. [Atlassian also reports](#) that 85 percent of knowledge workers now use AI, but only 29 percent have actually changed how they work. That gap, access without redesigned work, is exactly where planning has to do the heavy lifting. It will not close itself.

Contractors and AI agents are already absorbing headcount work

A real share of 2027 work will not be done by an employee. Contractors, outsourced functions, and AI agents are absorbing work that used to become a req. Companies that cannot control their own headcount today have little chance of controlling a mixed human and AI workforce in 2027. The control layer has to exist for people before it can extend to agents, which is the practical argument for fixing this now rather than after the plan locks.

THE DECISION IN FRONT OF YOU

The choice is not which planning tool to use. It is whether your two largest, fastest growing expense lines, people and AI, are managed by one system, or reassembled by hand every month in two different places. Load and lock is week 12 of the sequence in Part 8, not something to figure out after the plan is already approved.

A hire, an AI upgrade, or a digital worker now compete for the same dollar

Planning conversations still start the way they have for a decade: what do we want to spend, and on how many people. That question is no longer complete. A hiring plan used to have one shape: hire, or do not hire. It now has three shapes, and all three come out of the same budget.

- 1 **Fund a new hire**, who arrives with a salary and, increasingly, a real AI cost attached.
- 2 **Give more AI capability** to someone already on the team, raising what they cost without raising headcount.
- 3 **Fund a digital worker** that does the work instead, with its own budget and, in most companies, no approval process built for it yet.

KPMG frames the same underlying choice as *buy, build, borrow, or bot*: hire talent, develop it internally, use contingent labor, or use AI. Whatever you call the four options, the discipline is the same: price them consistently, or someone will guess.

The third option needs a name before you can budget it

An agent is a technical thing. It is an API key, a model, a registration in some system. You cannot budget an API key against a role for the same reason you cannot budget a laptop against a role. It is the wrong unit.

What you can budget is a digital worker: capacity you fund to do defined work, with a charter, a spend envelope, a cost center, and a human who owns it. One digital worker may run on several agents, models, and API keys, and those can change underneath it without changing the plan.

This is the same invention that made headcount plannable in the first place. A position is not a person. It is a funded seat that exists before anyone fills it, carrying a level, a budget, and a reporting line, which is why you can plan hiring a year before anyone is hired. The digital worker is the position object for the second workforce, and without it, your third option has nothing for a dollar figure to attach to.

Finance approves the hire. Someone in engineering or IT approves the AI seats. Nobody approves the digital worker, because there is no form for that yet. The 2027 planning problem is not forecast accuracy, but whether your plan can price a hire, a capability upgrade, and a digital worker against the same dollar and let someone choose.

RUN THIS ON ONE PERSON

Take your heaviest AI user in engineering. Pull twelve months of their actual tool and API spend and add it to their base pay. That total is what the role costs you. It is not the number sitting in your plan. That gap does not show up as a line item. It shows up eight months from now, in a cost center nobody was watching, and gets explained after the fact instead of planned for.

Not every AI dollar is a digital worker

Copilot licenses and developer tooling are software spend, and they should stay software spend. The test is substitution. If the spend funds work that would otherwise be a role, it belongs in the workforce plan. If it makes existing work faster, it belongs in the software budget and shows up in your plan inside a person's fully loaded cost.

Run that test once, deliberately, before departments start building. Companies that skip it end up with an IT budget quietly absorbing labor decisions, or a headcount plan cluttered with seat licenses that were never a hiring tradeoff.

FINANCE

You are being asked to defend a headcount number that no longer contains the full cost of the people in it, and a growing AI number that lives outside the plan entirely.

TALENT

You are being asked to hire against a budget that may already be splitting toward tools and digital workers in ways nobody has told you about.

HR

You own access policy, but AI tool access is spreading through expense reports and one-off purchases faster than any policy can track.

Five ways the 2027 plan breaks

The first four are the ones that have always broken a headcount plan. The fifth is new, and it is growing faster than the other four combined.

1. Drift: roles get opened that nobody approved

A hiring manager has leftover budget and a recruiter has capacity, so a req gets posted with no gate between intent and action. The org finds out when an offer needs signing, or when Finance closes the quarter.

"There's an enormous backlog of headcount that no one had ever accounted for. That's ridiculous."

A Talent leader at an advanced manufacturing company

The good news: a TeamOhana customer does not run into this blind. [Docker](#) stopped six figures of unapproved monthly spend once every request had to map to an approved position, and [SeatGeek](#) prevented 200,000 dollars of overspend in a single cycle.

2. The reconciliation tax: Finance rebuilds the forecast by hand

Every month, someone exports the HRIS, exports the ATS, opens the plan, and manually rebuilds the number so all three agree. Mid market Finance and Talent teams burn 60 to 120 hours a month on this. [Greenhouse](#) cut that to roughly two hours a week. [Metronet](#) saw a 7x return, the equivalent of at least three full time roles reclaimed from reconciliation work.

3. The approval black hole: no trail, no defensibility

A role was approved, probably. The evidence is a memory and a thumbs up in a Slack channel.

"Finance ends up actually manually rebuilding the whole headcount forecast, and you're pointing back to a thumbs up in a Slack channel."

A Finance leader at a life sciences company

This is a slow cost until it is a sudden one. It shows up hardest when a board asks who approved a position, or when a company heads toward an audit or a financing event. The

good news: [SeatGeek](#) replaced exactly this kind of operational chaos with one system of record Finance and Talent both trust.

4. Timing blindness: the variance nobody models

The plan is on track in headcount and off track in dollars, because a handful of start dates slipped a quarter and nobody repriced the plan when they did. [IonQ](#) identified 594,000 dollars in annual savings once timing was visible in dollars rather than headcount.

5. The token blind spot: AI spend nobody can attribute

This failure mode did not exist three years ago. It is now often the fastest growing number in the company, and it shows up the same six ways almost everywhere we have looked closely.

- Finance owns the AI number by hand, in a spreadsheet reconciled weeks after the money is already gone.
- The heaviest individual spenders are invisible, because their usage sits inside an enterprise plan while tracking only covers the API console.
- Contractor and intern API keys map to nobody, because contingent workers are frequently missing from the headcount system entirely.
- Autonomous agent spend hides on a cloud bill nobody in Finance owns or sees.
- The same dollars get counted twice across two different reporting sources.
- There is no baseline for normal, so a healthy engineering team's heavy usage and one runaway account look identical in a raw total.

[Deloitte](#) treats the token as the basic unit of AI cost and value, and recommends real time monitoring and forecasting discipline, the same rigor Finance already applies to headcount. [McKinsey](#) goes further: the unit that should be governed is the completed business outcome, not the raw token count. The good news: this is exactly what [Token Spend Management](#) is built to close, every AI dollar attributed to a person, team, or agent, inside the same system that holds the headcount plan.

SEE IT BEFORE IT HITS YOUR PLAN

Find out what the token blind spot is costing your team.

Twenty minutes, your own data. See fully loaded cost, comp plus AI spend, on one line.

[See a demo of Token Spend Management →](#)

FINANCE

Drift and the reconciliation tax cost you hours. The token blind spot costs you a forecast, since AI spend is a people cost wearing an IT invoice, scaling with who you hire, not with a software budget you control.

TALENT

The approval black hole and timing blindness are the ones that make you look unprepared in a leadership meeting, holding a plan you cannot fully defend.

HR

The token blind spot is yours to help fix, not because you caused it, but because attributing any AI dollar starts with having every contractor and intern in the headcount system with a consistent worker type. Most companies do not.

What planning season actually costs

To size these in your own numbers: count unapproved reqs from last year, tally reconciliation hours across Finance and Talent Ops, time how long it takes to answer who approved a role, compare planned to actual start dates, and compare per-head AI spend to what is currently budgeted. The table below separates the general benchmark from what a specific customer did about it, since those are two different kinds of evidence.

FAILURE MODE	INDUSTRY BENCHMARK	HOW A CUSTOMER SOLVED IT
Drift	Unapproved spend commonly reaches six figures a month before a gate exists.	Docker stopped six figures of unapproved monthly spend. SeatGeek prevented \$200K of overspend in one cycle.
Reconciliation tax	60 to 120 hours a month at mid market organizations.	Greenhouse cut it to about two hours a week. Metronet saw 7x ROI on the same problem.
Approval gap	Two to four week approval cycles are common without a system of record.	SeatGeek ended operational chaos and scaled without adding headcount to manage it.
Timing drift	Timing and level drift often outweigh headcount misses as a variance driver.	IonQ identified \$594,000 in annual savings once timing was priced in dollars.
Token blind spot	AI-native app spend grew 400% year over year at large enterprises.	Token Spend Management attributes every AI dollar to a person, team, or agent, inside the headcount plan.

The line that should worry Finance most: business units, not IT, now control 81 percent of software spend, against IT's 15 percent. AI is landing on someone's expense report inside a business unit before any central function sees it. By the time it reaches a budget review, it is already a run rate, not a decision.

How do you manage AI token spend alongside your headcount budget?

Annual planning has always been about deciding what to fund. What changed is that funding now splits three ways, and all three draw from the same pool of money.

OPTION	WHAT IT DOES	WHO USUALLY DECIDES TODAY	WHERE IT IS APPROVED TODAY
Fund a hire	Adds a person and, increasingly, a real AI cost attached to that person	Hiring manager, Finance	The headcount plan
Add capability	Raises how much an existing person can do, without adding headcount	Whoever owns the tool budget	Usually nowhere, or a procurement form
Deploy a digital worker	Produces output with no one at the keyboard, consuming its own budget	Engineering, IT, or nobody	Usually nowhere at all

That last column is the whole problem. One of the three options runs through an approval workflow. The other two do not. If your planning process cannot price all three against the same budget, your leaders are not making a tradeoff between them. They are guessing, and Finance reconciles the guess after the money is spent.

The data already shows it. [Gartner's CFO survey](#) found headcount growth expectations falling from 6 percent to 2 percent between 2025 and 2026, a pivot their analysts describe as structural. [KPMG makes the same point](#) from the workforce planning side: the useful classification is not "will AI take this job," it is whether a task is human led, AI assisted, AI automated with oversight, or not appropriate for AI at all. That classification is worth running task by task inside the Build phase of your own sequence, in Part 8, rather than deciding it once at the role level.

FINANCE You should be able to model all three options against one envelope before anyone commits, the same way you already model hire versus backfill.	TALENT You need to know when a role you are about to open is actually being weighed against a digital worker, so you are not staffing up to compete with a decision that has not been finalized.	HR You own the policy question underneath this: what work requires a human, what does not, and who decides. That is a people question before it is a technology question.
--	--	---

TeamOhana Computer: your AI coworker for planning

A chatbot waits to be asked. TeamOhana Computer connects directly to your ATS, HRIS, and finance data, already inside your access controls, and reasons through multi-step planning requests instead of one-shotting them. Ask it a question, drop in your own file, or assign it a task, and it shows its work, asks when it is unsure, and remembers your corrections.

Why a generic AI agent is not the same tool

Most teams have already tried wiring workforce data into a general purpose AI agent. It usually works, for one person. One team built their own agent so capable that, in their own words, only one person could use it, it was too powerful. A raw agent does not know your org chart, your approval chain, or who is allowed to see what, so the choice becomes either lock it away or hand out access nobody can safely audit.

	GENERIC AI AGENT + MCP	TEAMOHANA COMPUTER
Connects to your systems	Yes, if you build and maintain the connectors yourself	Yes, already built and normalized
Understands your data model	No, sees raw fields from each system separately	Yes, one reconciled layer across ATS, HRIS, and finance
Knows who can see what	No, access is usually all or nothing	Yes, permission aware by role, no back door to salaries
Stays current	Depends on whoever maintains the sync jobs	Near real time, always
Safe to hand to every hiring manager	Usually not, one power user holds all the access	Yes, by design

[Atlassian's own research](#) is a useful check on this: 85 percent of knowledge workers already use AI, but only 29 percent have changed how they actually work. Access to a chat window is not the same as a coworker embedded in the workflow with permission to act. That gap is exactly what a generic agent leaves open and what TeamOhana Computer is built to close.

The trust model

For any action that changes a real record, the agent runs a dry run first and asks for confirmation before anything actually changes. Nothing updates by accident, and every action

leaves a record of who, or what, made the change. Teams hand it real work because of that dry run, not only questions.

What Finance, Talent, HR, and executives are already asking it to do

FINANCE / FP&A

"I need to find \$1M in savings within our hiring plan with the least impact to our goals. Model these scenarios and rank them by savings versus business impact."

"Generate a dashboard that compares my May 1 headcount forecast snapshot to my July 1 forecast snapshot, and include recommendations."

TALENT ACQUISITION / RECRUITING OPS

"Show me all open headcount with a target start date in the next 90 days that don't have a recruiter assigned."

HR / PEOPLE OPERATIONS

"Generate an org design dashboard I can filter by division that assesses span of control, org layers, and key ratios."

EXECUTIVES

"Build a workforce health dashboard including average tenure, year to date attrition, headcount growth versus the same period last year, and average comp ratio."

See TeamOhana Computer on your own headcount data.

Try the scenario modeling use case first. [See TeamOhana Computer → Book a demo →](#)

Introducing the ledger for AI spend

Token data shows what was spent, but only the employee record explains who spent it and why. Most companies hold those two facts in different systems, so AI spend gets reported weeks late instead of planned in advance.

The idea worth taking into 2027 planning, regardless of which tool gets you there: one ledger for both. Every AI source, from Anthropic and OpenAI to Cursor and AWS Bedrock to seat based tools, consolidated with your HRIS and ATS data in one place. Pairing them makes five things possible.

- Every dollar attributed to a person, a team, or an agent, with an owner to contact when something spikes.
- No double counting, so when two sources report the same spend, you keep one clean number instead of two.
- Outliers measured against a team's own baseline, not a raw total, so one person spending ten times their team's average surfaces in days instead of at year end.
- Fully loaded cost per employee: compensation plus AI spend on one line, rolled up by team and by org.
- Forecasting at the point of the headcount request itself, so a planned hire carries expected token spend alongside comp, budgeted before the offer goes out, not discovered after.

Visibility and planning are not the same thing. [McKinsey argues the unit that matters is cost per completed business outcome](#), not cost per model call. A report tells you what already happened. A plan has to price what has not happened yet.

"The big thing is having all token spend in one place, so we can actually analyze what's going on. A new engineer next year adds incremental payroll and incremental AI spend. Budgeting for that is exactly where we're going."

Elisa Folden, Sr. Manager, Business Systems, SeatGeek

The one prerequisite that determines whether any of this works

None of the above works if your contractors and interns are not in your headcount system with a consistent worker type. They are typically the heaviest unattributed AI spenders and the population most likely to be missing entirely, because a sync filter written years ago to exclude consultants often catches fixed term and intern workers by the same rule. You cannot

attribute a dollar to someone your systems never loaded. Fixing this is the actual first step, before any ledger or dashboard.

See fully loaded cost on your own numbers.

Comp plus AI spend, one line, before the offer goes out. [See a demo of Token Spend Management →](#)

FINANCE

This is the missing line in your plan right now. Build the per-role baseline from your own data, three to six months of actual usage by role, then attach it to 2027 reqs the way you already attach comp.

TALENT

A hire's fully loaded cost now includes what they will spend on AI tools. That number belongs in the offer conversation with Finance before the req opens, not discovered after the person starts.

HR

Contractor and intern data quality is the dependency. Before planning starts, confirm every worker type in your headcount system is consistent, or the AI attribution work downstream is built on a gap.

The planning sequence

Twelve weeks, six phases, three owners, with an AI checkpoint built into every phase. Twelve weeks is the pace shown here; a smaller company can compress it to six to eight weeks, a larger one may need sixteen. The order and the ownership matter more than the exact calendar. On a calendar fiscal year, a 12 week sequence starting in early September has you board ready by the end of November. On a non-calendar fiscal year, count back twelve weeks from your board date and the sequence holds.

The critical design choice is in the last phase. Most companies stop at approval. This sequence ends at load and lock, where the approved plan, including its AI spend, becomes the operating constraint in the systems where work actually happens.

PHASE	WKS	OWNERS	AI CHECKPOINT	EXIT CRITERIA
1. Baseline	1-2	Actuals, open reqs, active roster reconciled	Confirm contractors and interns are in the headcount system with a consistent worker type	One agreed headcount and dollar figure, no side spreadsheets
2. Envelope	3-4	Total dollar envelope by department and quarter	Attach a per-role AI spend baseline to the envelope	Envelope already includes expected AI cost, not a body count
3. Build	5-7	Cost per role by start date; org design	Ask whether each role could be a digital worker instead	Every role has a level, a cost including AI spend, a start date, a reason
4. Pressure test	8-9	Downside and upside cases	Model AI capability added instead of new hires in the downside	Ranked cut/add lists include the tools versus hire tradeoff
5. Approve	10-11	Board and exec approval	Digital worker deployment on the same approval as a hire	Every position and digital worker has an approver, timestamp, cost
6. Load and lock	12	Budget loaded; variance thresholds set	AI baselines and digital worker budgets live in the same system	The approved plan is the operating constraint, not a document

A TOOL FOR PHASE 3: KNOW THE NUMBER BEFORE YOU WRITE IT DOWN

Pricing a role by level is only as good as the comp data behind it. TeamOhana's benchmark is built from more than 8,600 real, signed offers, seed to pre-IPO, not surveys. Enter a role and level and get the real range back in seconds at

Three rules that make it work

- 1 Give departments dollars, not headcount.** A body count invites level inflation. A dollar envelope, inclusive of expected AI spend, forces the requester to make the tradeoff themselves.
- 2 Never let AI spend live in a separate system from the headcount plan.** If Finance is reconciling two forecasts by hand, the plan is already out of date the day it is approved.
- 3 Never end at approval.** If phase six does not happen, phases one through five were a forecasting exercise, and next spring will prove it.

What happens after the plan locks

Mapped against the full lifecycle, the guide so far covers define, budget, approve, and meter. Review, change, and retirement are still missing, and retirement is the one that decides whether an AI budget compounds or ratchets.

A hire has a life after the offer. Onboarding, payroll, a review cycle, a promotion or a transfer, and eventually an exit. Every one of those events has an owner, a form, and a system, and every one of them updates what the person costs you.

A digital worker approved in your 2027 plan usually has none of that. It gets deployed, and then it runs. The only event anyone notices is the invoice, and the invoice arrives after the money is gone. Three things have to happen after approval, and they are the difference between a plan and a document.

Review it against its charter

You defined what the digital worker was funded to do. Review it against that, on the same cadence you review anything else that costs six figures. The question is not how many tokens it consumed. It is whether the work it was funded to do got done, and what it cost to do it.

This is also the honest answer to a measurement problem most companies are stuck on. Counting hours saved does not work, because saved hours get absorbed by other work and never show up in the P&L. What does show up is the budget. A digital worker with a charter and an envelope is measurable the same way a funded role is measurable, and that is the only measurement your CFO can act on.

Change it deliberately

Digital workers grow. Scope widens, a companion tool starts operating on its own, spend climbs past what anyone approved. Every one of those is a change to a funded position, and it should route through the same approval as a comp change or a level change. The alternative is autonomy expanding by drift, which is how a nine month old deployment ends up doing work nobody authorized at a cost nobody approved.

Decommission it and take the money back

Turning an agent off is a keystroke. Recovering its budget into the plan, where it can fund the next worker, human or digital, requires a decision and a system that holds the plan.

This is the step almost everyone will skip in 2027, and it is the one that compounds. A dollar you do not reclaim becomes permanent. Do this for two cycles and your envelope is

structurally larger every year for capacity you stopped using, and nobody can point to when it happened.

FINANCE Budget recovery at decommission is the difference between an AI line that compounds and one that ratchets. Decide now who is accountable for reclaiming it, and on what cadence.	TALENT When a digital worker is retired, the work it was doing does not disappear. That is a staffing question landing on you, often with no notice, unless retirement is a planned decision rather than a quiet one.	HR Every digital worker needs a named human owner, the same way every position has a manager. Ownership is what makes review and retirement possible, and it is the first thing to get lost when deployment happens outside a plan.
---	--	--

Where this sits in your stack

Your HRIS records human workers. Vendors are now shipping registries that record agents: identity, credentials, permissions, and an audit trail of what is running. Workday's Agent System of Record became generally available this year and does exactly that.

Registries answer a real question, which is what is running and who has access to what. They do not answer whether it should be running, what it is worth, or what it cost against the alternative you did not choose. Those are planning questions, and they get settled before anything gets registered.

Keep the two straight as you evaluate tools this season. A registry is where a digital worker gets provisioned and deactivated. A plan is where it gets defined, funded, approved, measured, and retired. You need both. Only one of them is where the money gets decided.

The readiness score

Sixteen questions. Check each box only if the answer comes from a system in under five minutes, and Finance, Talent, and HR would all give the same answer. Score the two sections separately, headcount out of 8 and AI out of 8, since almost everyone is strong on one and weak on the other. Run this out loud, together, before you build a single line of the plan. Disagreement about the answers is the most useful thing that happens in the meeting.

HEADCOUNT AUTHORIZATION

- ☐ You can state today's approved headcount and today's actual headcount, and the gap between them, without opening a spreadsheet.
- ☐ Every open req maps to a specific approved and funded position.
- ☐ A req cannot be opened without that mapping existing first.
- ☐ For any position, you can name the approver, the date, the approved level, and the approved cost.
- ☐ Your plan is expressed in dollars by start date, not only headcount by quarter, and reprices automatically when a start date moves.
- ☐ Every backfill required an explicit decision to refill, not an automatic one.
- ☐ Finance, Talent, and HR pull the same headcount number from the same place.
- ☐ You could hand an auditor a complete approval trail for your largest expense line without a project to assemble it.

AI AND AGENT READINESS

- ☐ Every contractor and intern who holds an AI tool account exists in your headcount system with a consistent worker type.
- ☐ You can state fully loaded cost per hire, comp plus expected AI spend, for any role in the plan.
- ☐ Every planned hire carries an AI spend estimate alongside comp, before the offer goes out.
- ☐ When a leader trades a hire for a digital worker, that decision is recorded somewhere, not just discussed.
- ☐ Your AI spend by person or team is measured against a baseline, not just a raw total.
- ☐

You know today whether an autonomous agent is running on a cloud bill nobody in Finance can see.

- ☐ Every digital worker in your plan has a named human owner and a written scope of work.
- ☐ When a digital worker is retired, you know how its budget returns to the plan and who decides.

___ / 8
HEADCOUNT

___ / 8
AI & AGENT

___ / 16
TOTAL

14-16 — You are running on authorization, for people and for AI. Move up a level: treat agent capacity as a planned resource with its own scenario modeling.

9-13 — Headcount is mostly under control. AI spend is probably still running on trust. Close the contractor and intern data gap first.

5-8 — You are running headcount on reconciliation and AI spend on nothing at all. Fix the authorization layer before you plan into it.

0-4 — Your two largest, fastest growing expense lines are both effectively uncontrolled between approvals. Treat this as a controls issue, not a tooling preference.

Score under 9? Let's talk through where the gap is.

Run the score with Finance, Talent, and HR in the room, then bring us the answers.

Book a demo →

What changes when this is one system

RUNNING ON RECONCILIATION		RUNNING ON AUTHORIZATION
The headcount number	Produced monthly by an analyst from four exports	Held continuously by one system all three teams read
The AI spend number	Assembled by hand from vendor portals, weeks after the money left	Consolidated daily, attributed to a person, team, or agent
A new hire's true cost	Comp only, AI spend discovered later	Comp plus expected AI spend, one number, before the offer goes out
A hire versus digital worker decision	Made informally, in a hallway or a Slack thread	Priced against the same budget, with a named owner either way
A slipped start date	Found at quarter close	Reprices the forecast the day it changes
An AI spend outlier	Invisible in a raw monthly total	Flagged against the team's own baseline within days
Finance in the process	Rebuilds two forecasts by hand and arrives late to both decisions	Sets one envelope covering both, and is present at the moment of decision
Audit readiness	A project, for headcount and for AI spend both	A query

What this looks like in practice

These are outcomes from companies that brought headcount and AI spend into one system.

Scale AI

Built a headcount planning engine to 10x its hiring velocity.

IonQ

594,000 dollars in annual savings identified through timing and level visibility.

SeatGeek

Ended operational chaos, 23 hours a week saved.

Docker

Six figures of unapproved monthly spend stopped.

Metronet

7x ROI, the equivalent of at least three full time roles saved.

Greenhouse

Reconciliation cut to roughly two hours a week, three teams aligned.

Rad AI

Scaled to 220 employees on a single source of truth.

Cedar

Freed up 1.5 million dollars by turning backfills into decisions.

THE HONEST CAVEAT

The AI spend and token attribution examples in this document are drawn from early customer work, current as of this planning season, not the multi-year track record behind the headcount numbers above. That is exactly why the readiness score in Part 11 scores headcount and AI spend separately. Most companies are further along on one than the other, and 2027 is when that gap needs to close.

Next step

Run the readiness score in Part 11 with Finance, Talent, and HR in the same room. If you score under 9 overall, the gap is not in your plan. It is between approval and execution, for people or for AI spend or both, and that is where 2027's variance will come from.

When you want to see this against your own data instead of a template, each of these is a short conversation away.

Salary benchmarks

Real comp data from 8,600+ signed offers, at benchmark.teamohana.com.

TeamOhana Computer

An AI coworker for headcount planning, connected to your live data.

Token Spend Management

The ledger that puts AI spend inside the headcount plan.

Position Management

The foundational layer this all sits on: approvals, backfills, and org data in one record.

Workforce Planning Platform

The full picture: headcount, compensation, and scenario planning in one system.

Planning season is weeks away.

Bring us your readiness score and we will walk your own numbers, twenty minutes, no template.

Book a demo at teamohana.com/book-demo →