

ai/r

Essay 1

Intelligence as Infrastructure

Redesigning Organizations for the Age of Abundant Intelligence

—

Todd Jochem

Principal Architect, ai/r

February 2026

The Light Bulb and the Grid

My parents remember when the Rural Electrification Authority first laid electric lines down the dirt road in front of their farm in southern Indiana. The very first use of electricity they had was a single light bulb in their house.

I think about that image a lot right now. Because that is almost exactly where we are with artificial intelligence inside organizations.

The groundwork is being laid — like those electric poles going up along dirt roads. And across the countryside of the modern enterprise, individual people are switching on their own light bulbs. Someone uses ChatGPT to draft a proposal. Someone else gets Claude to build a spreadsheet. A sales rep asks an AI to summarize a call. A finance analyst automates an invoice process. Light bulbs, flashing on and off, scattered across the organization.

And that's fine. That's how it starts. But a light bulb is not a grid.

The goal — the real transformation — is to make intelligence as ubiquitous as electricity eventually became. Not a novelty. Not a tool you pick up for a specific task. Infrastructure. Always on, always available, drawn on as needed, improving the quality of everything it touches. Electricity didn't just give us light. It reshaped how we built, worked, manufactured, communicated, and lived. Intelligence will do the same. We're just very, very early.

That shift — from scattered light bulbs to the grid — is what this essay is about. And it is the central argument of ai/r.

—

From Scarcity to Abundance

When I started in robotics and autonomy, there really was no such thing as artificial intelligence inside organizations. In fact, that's a very recent phenomenon.

Intelligence meant people — specifically, the people who had been there the longest, who had the most experience. That's where organizational intelligence lived. And if they left, you were in trouble.

Decision-making was structured hierarchically. The biggest decisions were made by the people furthest up the food chain, and the assumption was that their subject-matter expertise, their experience, their age, and their wisdom would produce good outcomes. Knowledge was scarce because the only real way to acquire it was through experience. There was no mechanism to convey institutional intelligence to younger people — they had their own jobs to do, and the only path to learning was doing those jobs, year after year. Senior management occasionally tried to lay things out, but that's very hard. There's no really good way to transfer decades of accumulated judgment.

That world is gone.

What's no longer scarce is the general ability to understand how things operate and what a reasonable response looks like. Computing is no longer scarce. And the fundamental difference — the thing that changes everything — is that artificial intelligence is now competent. Not perfect. Competent. Twenty years ago it wasn't. Five years ago it wasn't. Two years ago it arguably wasn't. It is now. It can work alongside humans. It can help. It can reason through problems. It can contribute meaningfully to how organizations operate.

This is a shift from scarcity to abundance. And most organizations have no idea what to do with it.

They act as if intelligence is still scarce — because large organizations are massive entities with inertia, and it's scary to change. Especially for public companies: your current processes generally lead to profits, and if they don't, you know which strings to pull. Adding something as disruptive as AI into that system is a big change, and big organizations usually don't like big changes.

And so you get friction. Everywhere. At every level. Senior managers know something has to change to remain competitive in the longer term. Middle managers are trying to figure out how they can still be useful — because they see AI as a direct threat to them, and in some ways they're right. And young people are scared. They don't know if what they just spent four, five, six years preparing for is about to become irrelevant.

I don't think it's quite right to call this a mistake. Many organizations are starting to embrace AI. But fundamentally, it needs to be embraced as an opportunity, not as a threat. And right now, in whole, it's thought of as a threat — because it's an unknown. I think it's a very big opportunity. For organizations and for people to redefine themselves. If they're willing to do that.

—

Tools Versus Infrastructure

Using AI tools means using intelligence to do specific tasks. Help me edit a document. Create a pivot table. Automate an invoice to the logistics supply chain. Summarize a report. That's the light bulb. Useful. Bright. Limited.

Redesigning the organization around intelligence is fundamentally different. It starts at the base layer: examining what your organization actually does — what is your core benefit to your customers? — and then moving up through the entire structure, embedding intelligence into how you provide that value. Some of that is process-oriented: faster invoicing, better scheduling, more efficient production flows. But much of it is deeper. How do you provide more value for the same price? How do you manage the disruption when employees leave and take organizational intelligence with them? How do you spread knowledge so it doesn't walk out the door with any single person?

The goal is for intelligence to become a true substrate — a shared foundation from the most junior employee to the C-suite — so that there is a common, shared understanding, a common intelligence, a common voice and process for how the

organization thinks, decides, and acts. That's something I think is lost on most people right now.

To put it plainly: intelligence as infrastructure means it's not something you add on. It's something that's there all the time. It helps guide what you want the organization to be, from the top to the bottom. From C-suite projections and growth strategy, to hard topics like right-sizing the organization or geographic distribution, to middle management resource optimization, all the way down to the shop floor — how do you minimize occupational hazards? How do you give employees the right tools to be most productive? How do you take what one person learned on Tuesday and make it available to everyone by Wednesday?

That's infrastructure. Not a chatbot. Not an app. A substrate.

—

The Intelligence Substrate

Here is the model I keep coming back to.

Organizations should not treat AI as a set of job roles to be filled or discrete tools to be deployed. Intelligence should become a shared substrate — a foundational layer that the entire organization sits on, from which capabilities can be drawn as needed.

Think of the shift from generators to the grid. Early factories each had their own power source. A single engine, driving a single shaft, powering a single set of machines. Then electricity became infrastructure — invisible, everywhere, drawn on demand. You didn't think about where the power came from. You just plugged in.

Intelligence is following the same arc.

From this substrate, specialized capabilities emerge when needed. Some persist for ongoing processes — a continuously updated demand forecast, for example, or a real-time supply chain monitor. Some solve a specific problem and dissolve. But the solutions remain documented and recallable. The intelligence can be summoned again, with full context, the next time it's needed. Think of it as emergent

intelligence: not pre-built applications, but capabilities that arise from the substrate when the organization needs them.

Humans retain the critical roles in this model. They identify the issues. They frame the context. They summon the intelligence. They interpret the outputs. And they retain accountability. Over time, human roles evolve — but accountability remains human. That last part is non-negotiable in my view.

I think there's even something akin to Moore's Law operating here. As the underlying intelligence improves, the organization's margins can expand in proportion — until you hit the process ceiling, at which point it asymptotes until the software gets better or some other structural change happens. But here's the important thing: the revenue acceleration can continue even when you hit those process-oriented limits. Because the ability to generate new ideas, identify new markets, and create new products is essentially infinite. Margin expansion comes first. Revenue expansion follows. And the ceiling on the second is much, much higher than the ceiling on the first.

That compounding — intelligence improving the organization, which feeds better data back into the intelligence, which further improves the organization — is the flywheel. And once it's spinning, it becomes very difficult for competitors without a substrate to keep up.

—

What a Tuesday Morning Looks Like

Let me make this concrete. Imagine a 500-person manufacturing company that has adopted intelligence as infrastructure. What's different on a random Tuesday morning?

Across all levels, employees arrive and — through whatever mechanism, their phone, their laptop, a brief spoken summary — they're given real-time updates. The status of the company at a level appropriate for their role. What their day looks like. What they need to get done. And how they're going to be helped to get those things done.

The weekly executive meeting is short and sharp. The numbers are there, but so is something deeper: what the numbers actually mean in the short term, the long term, and the very long term. How that impacts staffing, inventory, purchasing, pricing, sales. Previously, that kind of analysis was done quarterly — maybe. Now the deeper insights are available nearly daily. There's noise in that daily signal, and you wouldn't want to react to every fluctuation. But all those deeper insights can be presented daily. And instead of working to produce those insights, the executives can spend their time thinking about them and deciding what to do next. That's a fundamental shift: from generating analysis to consuming it and acting on it.

The middle manager's workflow has shifted. They are no longer primarily supervising tasks. They're optimizing people and resources. I can easily see a world where junior employees are no longer narrowly specialized. They're high-agency people, multifaceted, trained across multiple domains, who understand they have deep support from the substrate to get their jobs done in all those areas. The manager's job is to orchestrate them — figure out where they're needed most on any given day or week, and make sure the intelligence is supporting them properly.

For new employees, this will just be natural. It will be how they work. There will be nothing different. The expectation is that they have agency, that they can solve problems, that they know they have tools and how to access those tools. For existing employees, the transition is harder, but the same principle applies: you're a jack of all trades now, and you have deeper support than any specialist ever had.

On the sales side, the substrate fundamentally changes how the loop between sales, production, and order management works. Intelligence can read the tone and frequency of emails from salespeople and their customers. It can make best estimates about what resources will be needed a week, a month, a year in the future. It starts optimizing for those — or at least letting people know what's coming down the pike and giving them options to plan based on their expert knowledge. It can predict what customers are going to ask, suggest responses, model how different customer types typically buy. Do they wait until the last minute? Do they want to lock things in early? What incentives have worked historically? What industry-wide trends can help

us differentiate? Salespeople aren't replaced. They're amplified. And the speed at which the organization responds to market signals goes from weeks to hours.

On the manufacturing floor, production scheduling is optimized against real constraints — what equipment is available, what orders are anticipated, what supply chain signals look like. There's only a certain amount of time you can shave off building a widget. But all the process and paperwork flow that never really needed a human anyway? That gets shaved off entirely.

And here's the part that really matters: when a disruption happens — a supplier has a fire, a region floods, a key customer goes on vacation and their orders are delayed — a human tells the substrate once, in one place, and the impact ripples everywhere it needs to go. Manufacturing adjusts timelines. Sales communicates delays to customers automatically. Purchasing identifies alternative suppliers and presents options: this one costs three percent more, but it arrives on time — what would you rather do?

Decisions that previously required days of escalation now take minutes. Not because humans are removed from the loop, but because the intelligence has already assembled the relevant information and presented options with context. Consider a major purchase decision: the substrate knows from seventeen of your twenty-seven sales managers that there's strong interest in a product. It knows historically what that level of interest translates into. It can analyze all of that, determine when orders are most likely to arrive, and make recommendations about inventory allocation and purchase timing. The decision doesn't arrive as a request to be researched. It arrives as a decision to be made, with the reasoning already laid out.

And some things genuinely don't require escalation anymore. If a manager has consistently authorized a ten-percent price break for certain customers, that's a decision people would be comfortable delegating to the substrate. Those are easy examples. But they add up. The speed and quality of decisions across the organization gets measurably higher.

That's what a Tuesday morning looks like. Not science fiction. Just intelligence woven into how the organization actually operates.

—

The Human Layer

None of this works if you skip the hard part. And the hard part is not technology. It's people.

AI threatens things that humans hold onto tightly: expertise status, management authority, executive control narratives, decision ownership, and institutional ego structures. These are not irrational fears. They are structural realities.

Senior leaders often know that something has to change, but the act of changing feels like surrendering control to something they don't fully understand. Middle managers see intelligence as a direct threat to their relevance — and in some ways, they're right. Much of traditional middle management involves aggregating information, summarizing it upward, and distributing directives downward. A substrate does that faster and more completely. Junior employees wonder if the skills they just spent years acquiring are about to be devalued.

Everybody's scared. And fear, unaddressed, produces resistance. Not rational objection — resistance. The kind that shows up as delay, passive non-adoption, political maneuvering, and a thousand small decisions to not engage.

This is why ai/r addresses the human layer first. Not because technology is unimportant, but because no architecture survives contact with an organization that hasn't reckoned with what intelligence integration actually requires of its people.

Who must be aligned before anything meaningful happens? I think it starts at the very bottom and at the very top — because the middle is the hardest. The middle is where it's messiest. The middle is where there's this confluence of intellectual capital, accumulated knowledge, and the chaos of process. At the top, you can think clearly

and strategically. At the bottom, simple things pay off immediately. The middle is where all the friction lives.

Getting the rank-and-file employees bought in is essential. Presumably, if this is happening, upper management has bought in and allocated the resources. But if the people doing the real work every day don't buy in, it doesn't work. And that's a human problem, not a technology problem.

I've learned something about this — from coaching teenagers, from building companies, from leading engineers. If you can find the right people, they'll want to succeed for you. They'll do what it takes. They'll work hard. They'll sacrifice for the good of the team, knowing that in the end it's best for them too. That's what coaching taught me about human adaptation. And I think this is the exact same case. We need to get the right teams together in these organizations and move forward.

What I'm not arguing is that everything will be fine. There will be disruption. There will be employees who need to move on and find new jobs, because they don't fit, or they don't want to work with it, or their jobs genuinely aren't needed anymore. And that's going to happen quickly. But I'm optimistic that we can find ways to handle that. I don't know exactly what those ways are yet. But I'm certain we can find them.

What I'd rather see — and what I think most organizations would rather see — is the opposite of displacement. If you have a group of high-agency people, five or six or ten, who can work together and produce the results of ten or twenty times as many people because they're leveraging the intelligence substrate? Let's do that. That's even better than reducing headcount. That's multiplication, not subtraction.

But if people eventually refuse to change — if they refuse to embrace a direction that's needed for the organization to survive in the long term — then those people have to move on. And that incentive is exactly the same as it's always been. The question is whether humans will make those hard decisions based on non-human input. We're giving humans something that's clearly threatening and requires them to change. Will they use it or not? Other humans have to assess that.

—

The Orchestrator Layer

Middle management doesn't disappear. It transforms.

The routine oversight shrinks. The orchestration role expands. Middle managers become integration stewards — human-machine orchestrators who ensure the substrate is being used well, that human judgment is being applied where it matters, and that the organization is learning from both its people and its intelligence.

My own experience with entrepreneurship taught me something relevant here. I ran small companies where it was me, a level of PhDs, and the workers. Extremely flat. The subject-matter domain experts at the middle level didn't have direct lines to the people doing the work in the traditional sense — instead, they had the freedom for their knowledge and intelligence to be shared across all parts of the organization. They had their fingers in multiple projects because you needed that in a small company.

I think that kind of structure could come back. Not that you necessarily need PhDs, but the shape is the same: flat, broad, with domain experts whose intelligence flows across the whole institution rather than down a single reporting line. Fewer direct reporting lines. More cross-functional influence. That's very similar to entrepreneurship — and I think it's where larger organizations are heading too, enabled by the substrate.

—

What Changes About Authority, Expertise, and Careers

If intelligence becomes truly abundant inside an organization, some uncomfortable questions follow. What happens to authority? What happens to expertise? What happens to career ladders?

I think intelligence practically flattens organizations, even if human nature keeps some hierarchy in place. Human nature says there will always be hierarchy. People still want leaders. Some people don't want to be leaders — they want to be the person

with their hands in all the pies, doing the work. That continues. But the organization gets flatter. Power distributes. And it should distribute — carefully, but it should.

Authority changes in a specific way. When intelligence is abundant, people in authority have to really think through their thinking. Personal agendas largely lose their hiding places. Everything is on the surface. You know the reasons why people in authority make the decisions they make. And maybe, counterintuitively, trust deepens — because the basis for decisions becomes transparent.

Tenure may lose some of its traditional power. That's a real possibility. But the role of good leaders has remained essential for thousands and thousands of years, across technology revolutions far more dramatic than this one. Good leaders are good leaders. You're going to need them. But in a flatter organization, younger people could move to leadership roles more quickly and have a much bigger impact. That's not a bad thing.

Expertise shifts meaningfully. The ability to think across multiple domains matters a lot more. Specific, narrow expertise matters a lot less. We're already seeing this in software: if you can design the system, you're much more valuable than if you can write the code. I think that's a good analogy across every kind of work. If you can understand the market, you're more valuable than if you can sell one niche. Breadth of thinking becomes the premium skill.

And career ladders? This might be the most interesting change of all. Maybe instead of measuring a career by the height you climb, it's measured by the breadth you reach and the people you impact. If you're brilliant at understanding markets and people but don't want to run an organization, you'll have the ability to work across different organizations at the same level, doing meaningful work, being rewarded throughout your career. Lateral movement stops being a neutral or negative signal and starts being a sign of range and impact. Your ability to think broadly and communicate insights effectively to the substrate and the intelligence of whatever organization you're in — that's going to be incredibly valued.

Compensation logic has to evolve too. If an employee has the best ideas — ideas that improve the intelligence substrate across the entire organization — they should be rewarded for that. No different than being rewarded for putting a great suggestion in the suggestion box. Except now the suggestion box is connected to everything. Instead of having to move up the ladder to make more money or get more time off, maybe you can have just as big an impact staying where you are. And I think that's a fundamentally positive thing for high-agency people.

All of this changes who you hire, too. You want high-agency people who are comfortable in a co-working environment where they're relying on an intelligence substrate specific to your organization to assist with decision-making. But also — and this is important — people who are comfortable giving the intelligence what they've learned, so that others in the organization and the substrate itself can become smarter.

That requires a collegial mindset. Probably a stronger liberal arts orientation, even for high-technology applications. People who can communicate, synthesize, and think across domains. People persons, in the best sense.

Although, having said that, there's an interesting wrinkle. In many cases, AI does a better job relating to the people on the fringes — the exceptionally talented individuals who don't fit neatly into traditional team dynamics. So maybe there's more tolerance for different kinds of people in an intelligence-augmented organization, not less. The substrate can bridge communication gaps that humans struggle with.

—

Architecture Before Deployment

You have to understand how intelligence can work in your organization before you figure out how to put it in people's hands. That's the core principle: architecture before deployment.

This doesn't mean waiting forever. It means having a plan that starts with specific applications and expands outward deliberately. How do you want intelligence to be used? What are the limits? What are the institutional secrets you need to protect? Do you house it internally? Is some external? Does the model supplier give you adequate assurances that what you use it for is safe? There are no perfect answers to these questions, but you need to know what the answers are before you move forward at scale.

What happens when organizations skip architecture? That's where we are right now. Individual units and people are using AI on their own, and that's fine — it needs to happen that way. But that's fine when there are ten or fifteen or twenty people doing it. If there are ten thousand people all using it their own way, with their own tools and their own prompts and their own workarounds, it becomes unmanageable. Risks escalate. It'll be a challenge to manage risk and information across ten thousand people even with a unified substrate. Without one, you're not managing it at all.

The first mistake most companies will make? They'll delay too long and not think broadly enough about it. And then when they finally move, the opposite problem: the call for change and the sheer number of possible improvements hits like a tidal wave. Everyone wants everything at once. It becomes chaos, because it's never been done before. And potentially the underlying intelligence and hardware — the literal infrastructure that runs the intelligence — isn't ready yet because it all happens so fast.

If I walked into a two-thousand-person company tomorrow, day one would be simple: create a survey that asks every employee how they're already using AI and what they like and don't like about it. In the first ninety days, you roll out a handful of standardized AI capabilities — five or six well-chosen light bulbs — but with back-end hooks that allow everything to be connected and to start learning from each other.

I think of this in two phases. Phase one is **seeding the substrate** — getting intelligence into the ground, getting people using it, learning what works. You have to invest in the substrate before you get any value. Simple things first: helping people

write, organizing emails, drafting presentations. Those are boring. But they set the tone and evolve into something much bigger — a state where intelligence is everywhere in the organization, and people can converse with it, bounce ideas off it, and develop models of how things work using the assets and tools of the organization.

Phase two is **emergence** — when intelligence starts to surface across the organization in ways you didn't plan for. You know the shift is real at two moments. The first is when employees buy in and start actively showing you ways they've used AI that you didn't think of, to make their jobs easier. The second is when the true substrate nature of intelligence comes to play — when you realize large financial or market opportunities that you wouldn't have seen before, because now you have data that's integrated and analyzed across the entire organization. That's the flywheel spinning. That's the sweet spot you want to maintain.

—

The R in ai/r: When Intelligence Gets Physical

So far, I've been talking primarily about the cognitive substrate — software intelligence flowing through an organization. But ai/r has a slash in it for a reason. The R stands for robotics. And at some point, the intelligence substrate doesn't just think. It acts in the physical world.

Where does robotics integrate with substrate intelligence first? In companies that already have automation systems, it's easier. Think about logistics companies like UPS or FedEx that already use automation at scale. Those integrations happen first. And for most organizations, at least for a long time, robotics means one thing: logistics and physical labor. Moving things around. It's not sexy. It's not humanoid robots having conversations in the hallway. It's robotic systems moving stuff — in warehouses, on manufacturing floors, through supply chains. That's the first thing that gets integrated, and it's where the clearest value shows up.

I don't know that I see how physical embodied robots in an office environment are particularly beneficial, at least not yet. But in a manufacturing environment, they clearly are. And there too, it's the logistics side of manufacturing that leads — whether it's humanoids or otherwise, the initial value is in moving materials, sequencing operations, and optimizing physical workflows.

The substrate eventually controls physical systems, yes. But I want to be precise about what that means right now. Software already controls physical systems — that's not new. This is the next level up: the intelligence substrate optimizing when and how those physical systems run. I'm not talking about reprogramming machines on the fly. I'm talking about sequencing them, turning them on and off, scheduling them against real-time demand signals. Eventually it may get to deeper reprogramming, but that's not what I'm thinking of yet.

What's different about cyber-physical systems versus pure software? The stakes are higher. There's a much greater chance of physical harm to employees if something goes wrong. There's the fatigue factor — physical systems interact with humans who get tired, while software doesn't care about a third shift. And there's the real-world sensing dimension: one of the great things about using robotics in controlled environments is that the sensing can be very good. You control the environment, so the data quality is high. Human sensing is still as good as robot sensing for many things, and will be for a while. But the ability to sense reliably in these environments won't be an issue for most robotic systems.

The real power is in the feedback loop. The physical systems notice inefficiencies — or more precisely, the data they generate reveals inefficiencies to the substrate. Maybe you start noticing patterns in when trucks arrive, or how material flows bottleneck at certain points, or where the handoff between human and machine work creates friction. That data feeds back into the intelligence layer, which optimizes the physical operations, which generates better data, which further improves the intelligence. The same flywheel that operates in the cognitive substrate operates in the physical world too, just with higher capital costs and higher stakes.

And that's the capital reality: physical things cost more than software to implement at scale. So the capital largely gets spent on software first and slowly moves to physical systems. Eventually, in a mature organization, that ratio reverses — the physical infrastructure becomes the dominant investment. But that's at the middle and end of the adoption curve, not the beginning.

What could slow this down? Honestly, an accident. A serious one. These physical systems interact with humans and with the environment. You cannot have a major, catastrophic event. I'm almost talking about a black-swan scenario — something that makes headlines and sets the public narrative backward. I don't think that's likely. I think these systems will be much more self-contained than people fear. It'll be similar to how autonomous vehicles work now: people complain a lot, but the bottom line is they're not running around causing daily catastrophes. That's the bar. And I think robotic systems in organizational settings will clear it.

But it's a risk that has to be managed architecturally. You don't bolt robots onto an organization the same way you don't bolt software intelligence on. The physical layer needs the same structural thinking as the cognitive layer — maybe more, because the consequences of getting it wrong are tangible and immediate.

—

Where the Money Shows Up

In a thousand-person company, intelligence infrastructure moves EBITDA first in the simplest places. The one-off light bulbs making people more efficient. The same headcount doing more work. Productivity increase. It shows up in margin expansion before revenue acceleration.

G&A drops first — that's the lowest-hanging fruit. People can do more things on their own. As you grow, you don't need to hire proportionally. The G&A largely stays the same initially even as the organization grows, then drops through attrition, and eventually shrinks structurally as the technology improves.

Then, as the substrate matures and becomes truly ubiquitous, revenue expansion happens. New products, new markets, new capabilities that weren't possible before. You can even imagine scenarios where margins improve while prices drop — because the efficiency gains are that significant. Intelligence doesn't just help the organization. It can help the customer too.

Working capital allocation changes in an interesting way. Forecasting improves, so capital becomes more targeted. You allocate smaller sums more often. You try a lot of small things. And when you identify something that really works, you can go all in on it confidently — because you have more data, better models, and a clearer picture than you've ever had. There are fewer medium-level near misses. More small experiments and more big, confident bets.

By its nature, intelligence can generate tons of ideas. You need to be prudent about which ones you pursue — you can't afford to chase all of them. And some will fail. I don't believe intelligence is omniscient. It can't model the world at a level that predicts exactly what to do next. There will be things it gets wrong. But the difference is you go down those paths faster, which costs more money potentially, but also gets you to a solution faster. Those are the trade-offs you have to get used to.

How would I explain substrate ROI to a skeptical CFO? CFOs are often the stick in the mud here. My pitch is simple: this is the equivalent of adding PCs to your workflow thirty years ago. The equivalent of selling online twenty years ago. It's an investment you have to make. The good news is, if you do it right, this investment starts paying off far sooner than either of those did. Walk into almost any company today and you can already point to situations where AI is helping them be more productive. The question is whether you want a hundred scattered light bulbs or a grid.

The clearest financial signal of substrate maturity? Increasing margins as revenues increase — simultaneously. That's unusual. It won't happen early, because you're still installing and learning. And it may not persist forever, because competition catches up. But in the medium term, the possibility of growing both profitability and revenue

at the same time is real. And for companies that get there first, the advantage compounds.

—

The Private Equity Lens

Everything I've described so far — the substrate, the economics, the organizational transformation — applies to any company with meaningful process flows and knowledge work. But there's another dimension. Private equity firms don't just need to understand this transformation for a single organization. They need to evaluate it across an entire portfolio — and in every potential acquisition target they're considering. Which portfolio companies are best positioned for intelligence integration? Which targets have already started building a substrate, and which are falling behind? Where is the margin-expansion opportunity real, and where are the human and cultural barriers that would prevent it from materializing?

These are the kinds of structural questions PE firms need to be asking. And they need someone who understands real deployment, organizational systems, robotics realism, enterprise politics, and accountability — not another strategy deck from a big consultancy or another SaaS vendor pitch.

The opportunity for PE is to use the intelligence substrate as a lens for evaluating and transforming portfolio companies. Which companies in the portfolio are best positioned to benefit from intelligence integration? Where are the margin-expansion opportunities that automation can unlock? Where are the workflow inefficiencies that a substrate could address in the first ninety days? And just as importantly: where are the human and cultural barriers that would prevent integration from succeeding, even if the technology is ready?

These are structural questions, not software questions. And they require someone who can walk a manufacturing floor, sit in a board room, and understand both the robotics realities and the organizational politics.

There's a competitive dimension too. Established portfolio companies need to watch for smaller, faster competitors who embrace intelligence early. Smaller companies generally move quicker — they don't carry the inertia. A mid-market firm that builds a mature intelligence substrate could punch well above its weight class. PE firms that can identify which of their portfolio companies are most ready for that transformation — and which acquisition targets have already started building it — will have a significant advantage in the coming cycle.

Substrate maturity can become a genuine competitive moat. But the moat isn't the software. It's the intelligence built on top of it — the organizational knowledge drawn from employees, the accumulated decision-making patterns, the feedback loops between the physical and digital operations. The quicker a portfolio company can build that, the bigger its advantage. And the PE firm that understands how to assess and accelerate that process across a portfolio is operating at a level most of the market hasn't reached yet.

—

Governance and Accountability

Boards need to think about AI differently. Not as a technology initiative. As an intelligence initiative. What is our organizational intelligence, and how do we distribute it to all levels so employees can maximally do their jobs for the maximum benefit of the organization, themselves, and society?

The governance responsibility is the same as it's always been. Boards need to make sure that what they provide to employees — and eventually to customers — aligns with the corporate, social, and moral values the organization has developed. The mechanism for ensuring that changes. The responsibility does not.

Formal governance needs to kick in once you're beyond the light-bulb stage. People have individual agency when they're just using AI on their own. That doesn't change. But when intelligence starts affecting the entire organization and making organization-wide recommendations, governance structures need to be in place.

Who is accountable when the substrate recommends a course of action and it goes wrong? How is oversight structured? What's the escalation logic?

There's an important ethical dimension here too. Intelligence gives you the real opportunity to measure things that, from a pure dollars-and-cents perspective, might not seem to make sense — but from a macroscopic view matter enormously.

Employee retention. Customer loyalty. Community reputation. Organizational culture. These things have much bigger long-term financial implications than most quarterly reports suggest. Intelligence infrastructure lets you actually see and measure that, rather than guessing at it. You have to be careful not to make every decision about dollars and cents — but now you have the tools to show why the things that seem soft actually aren't.

And there's a counterintuitive upside. Intelligence infrastructure can serve as a form of organizational guardrail. Boards can use the substrate to enforce boundaries, to ensure decisions align with values, to create consistency in how principles are applied across the organization. It's accountability infrastructure — the same kind boards are supposed to provide, just more consistently and transparently applied. Boards exist for a reason. If you trust them, then they should be able to use these tools to protect shareholders, protect the organization, and protect the people.

Can intelligence infrastructure make bad leadership more dangerous? Yes. Like any powerful tool, it amplifies whatever it's pointed at. Bad intentions scale along with good ones. But so do the safeguards, if you design them in from the start. That's why architecture comes before deployment.

—

The Long View

What happens to organizations that don't redesign for abundant intelligence? I think in the short term, it's a slow decline. In the medium and long term — five to ten to twenty years — they shrink or disappear. Not every kind of organization, and not all at once. Not the grass-cutting company or necessarily the home builder. But for

typical consumer and industrial organizations with large process flows, this is existential. They can continue to find niches. But they'll shrink. And that's exactly identical to every other technological revolution of the past hundred years.

The biggest long-term risk is that you don't embrace this and someone else does. There will always be bad cases and bad uses — in every industry there have been for thousands of years, and there will continue to be. But you can't miss this boat. And the boat is coming much quicker than other boats have.

The biggest long-term upside? You can increase productivity, increase profitability, increase employee satisfaction, and decrease waste. It affects every single part of the organization. For employees and organizations that embrace it, it could really open up a redefinition of what working means.

I think ai/r — artificial intelligence and robotics, together — is coming faster than any previous revolution. The personal computer revolution took decades. The internet took ten years. This will take two, three, maybe four years to fundamentally reshape how organizations operate. That's both exciting and terrifying.

If someone reads this essay ten years from now, what would make me happy? If the idea of an intelligence substrate has become real. If the concept of emergent, on-demand intelligence — called upon to shape decisions, processes, and strategic direction — is simply how organizations work. That would mean we got it right.

—

Why I'm Writing This Now

This is almost a eureka moment for me. My background in robotics and self-driving cars. The way I was educated and trained to think in systems. A practical background of leading people in the companies I started and grew. A more emotional kind of leadership — developing winning teams and guiding young people through growth phases when I was a coach. All of it converges at this one moment in time.

I'm not an implementer. I don't want to install chatbots. I want to help people think about this problem and come up with novel solutions that help them, their organizations, and their employees grow.

I'm a natural systems thinker — I have been my whole life. I built self-driving cars, but those cars were systems: software that talked to sensors that created maps from those sensors that mapped different geometries onto those maps, and then finally made decisions based on what those integrated maps and geometries looked like. It was never one specific thing. It was always: how do I use technology to put things together to do something larger? It took thirty years, but it feels like we're in that same spot again — this time at the scale of entire organizations.

What do most AI commentators miss? I think they miss the systems layer. The big picture. Light bulbs are nice, but I care more about the grid underneath. Why am I thinking structurally instead of tactically? Because it's my nature. I've always built systems. And building systems means understanding how pieces connect and how the whole becomes greater than the sum of its parts.

Am I an eternal optimist? I'm a technology optimist grounded in realism. I've been around AI and robotics for thirty years. I'm not scared of it, but I'm cautious — because I've seen it grow and I understand it's not perfect. If you're expecting perfection, you'll be disappointed. That doesn't mean it can't be incredibly useful and transformative in nearly every field. And I think that's the kind of attitude you need when deploying these things. Not fear. Not blind faith. Informed confidence.

There will be disruption. There will be heartbreak and failure stories alongside the successes. The path will be hard. There will be hardship in short spurts, things that go wrong, people who are displaced before the system finds its balance. I'm not arguing that everything will be fine. I'm arguing that the outcome — for organizations, for societies, and specifically for people — will be better. And that the optimism is maintained even though the path is difficult.

—

What ai/r Becomes

Eventually, inside an organization that gets this right, intelligence becomes a coworker. It becomes part of the corporate identity. You could even envision it becoming part of the products you sell. It's a reflection of the people in the organization — their intelligence, their memories, how they communicate and solve problems.

That's the aspiration. Not that someone says, "We deployed AI." But that they say:

| *It's part of us.*

Integration at the identity level. Culture, systems, decision architecture — not tools bolted on.

If I had to compress this entire argument into three sentences: ai/r is coming. Ubiquity across your organization is going to be a necessity. And find people who want to go along on the ride with you.

If I had to compress it into one: ai/r can make us all better.

The light bulbs are already on. The question is whether you're building a grid.

—

Todd Jochem is the founder and Principal Architect of ai/r (air atelier), a boutique AI-native intelligence architecture studio. He brings early autonomy-era experience, real-world robotics deployment, and decades of organizational leadership to the question of how intelligence becomes infrastructure inside real organizations.

ai/r | airatelier.ai