

LEADERSHIP BRIEF

THE AI READINESS GAP

Why Most Businesses Will Get AI Wrong & What to Do Instead



“

*AI is the most efficient way to discover what's broken in your business.
Most companies make that discovery after the tool is already live.*

-Jeremy Kushner | CEO | BACS Consulting Group

”

THE **PROBLEM** NOBODY IS TALKING ABOUT



The AI story being told right now is almost entirely wrong.

It goes like this: AI is transforming business. The tools are powerful, accessible, affordable. Move fast or get left behind.

So thousands of businesses are following the script. Buying platforms. Signing contracts. Announcing AI initiatives in all-hands meetings. Checking the box.

And a lot of them are failing. They just don't know it yet.

Not because the technology doesn't work. It does. The models powering today's AI tools are genuinely powerful, and they will change how businesses operate. That part of the story is true.

The failure is happening somewhere else. In the gap between what AI can do and what most businesses are

actually prepared for.

“

Call it the AI Readiness Gap.

”

I see it regularly. The company that deployed an AI tool on top of data scattered across four systems that don't talk to each other. The business that mandated AI adoption without ever asking whether its people were ready, and watched the initiative stall the moment the initial enthusiasm wore off.





The organization that bought a sophisticated platform without a clear problem to solve and is now paying a monthly subscription for something nobody fully uses.

The AI Readiness Gap isn't a technology problem. It's a leadership problem. And it is far more common than the vendors selling AI solutions would like you to know.

And, like most leadership problems, it has a financial cost, in wasted spend, in failed initiatives, and in the competitive ground surrendered to the organizations that got it right.

“
AI does not replace operational discipline. It exposes the absence of it.
”

This paper is the conversation most of them won't have with you. What to do before you buy anything. Why that sequence matters more than which tool you choose. And what the businesses actually getting AI right are doing differently.

It will not tell you which platform to buy. There are plenty of people willing to do that.

What it will do is give you a clear picture of what AI readiness actually requires, and why most organizations are further from ready than they think.

THE FOUNDATION HAS TO COME FIRST

Before your business does anything with AI, one question is worth asking.

Not “which tool should we use?” Not “what’s the ROI?” Not even “are we ready?”

Simpler than all of those and harder: what is AI actually going to work with when we turn it on?

AI doesn’t arrive in your business and start fresh. It inherits what’s already there. Your data. Your systems. Your processes. Every workaround your team built because the official way stopped working three years ago. It takes all of that as its raw material.

And then it amplifies it.

Clean data, integrated systems, documented processes - AI amplifies those. Results come faster. The investment pays off.

Siloed data, broken workflows, information scattered across systems that don’t talk to each other - AI amplifies those too. Just as fast. What you get is a more expensive version of the problem you already had.

This is the central fact of AI adoption that most vendors would rather you not dwell on. I call it the amplification problem. And it is the most common reason implementations fail.

In practical terms: organizations with fragmented operations aren’t just buying a tool. They’re buying a faster, more expensive version of the problems they already have.



THREE PROBLEMS I HEAR IN ALMOST EVERY CONVERSATION



The first is siloed data. The CRM doesn't talk to the ERP. Financial data lives in spreadsheets because nobody trusted the actual system. Operational records maintained in ways that make sense to one team and nobody else. AI needs data it can work across. Siloed data gives it fragments.

The second is undocumented process. The way work actually gets done in most organizations is only partially written down. The rest lives in the institutional knowledge of the people who have been there long enough to figure it out. AI can't access that knowledge. It can only work with what exists in a form it can read. When the process isn't documented, the output reflects that.

The third is the trust problem, and it is the most telling of all. In many organizations nobody fully trusts the numbers. People run their own calculations. Maintain their own spreadsheets. Check the numbers themselves before they trust them. That isn't a data problem. It is a signal that the data foundation has never been resolved. AI doesn't fix it. It inherits it.

Most organizations already know these problems exist. They have been working around them for years. What AI changes is the cost of ignoring them.

WHAT GETTING THE FOUNDATION RIGHT ACTUALLY LOOKS LIKE



The businesses I see making real progress with AI didn't start with AI. They started earlier. A deliberate effort to consolidate data, integrate systems, document how work actually gets done.

Unglamorous work. The kind that never gets announced at an all-hands meeting. But it is what determines whether AI delivers on its promise or becomes a monthly subscription nobody fully uses.

The foundation doesn't have to be perfect before a pilot starts. A narrow, well-scoped application, clean inputs, a clear measure of success, can deliver real results even while the broader work is still in progress.

Think of it as Minimum Viable Data, not immaculate, not comprehensive, just clean enough for the specific problem you're solving.

The most practical approach is to let the pilot define the cleanup rather than front-load it. The first implementation surfaces exactly which data gaps matter for that application. Fix those. The cleanup becomes part of the work, not a prerequisite for it, and the pilot itself makes the business case for doing it.

The businesses that get into trouble are the ones that overestimate how ready they are. They discover the gap after the contract is signed.

“—

Know what you have before you decide what to build on it.

—”

THE HUMAN SIDE NOBODY PLANS FOR

Before almost every AI implementation I have been part of, the same conversation doesn't happen.

Not about data. Not about systems. Not about vendors.

The conversation about people.

Specifically: how employees feel about AI arriving in their workplace. Whether they see it as something happening with them or to them. That distinction determines more about whether an implementation succeeds than any technology decision you will make.

Most organizations never ask.



THE MANDATE THAT BACKFIRES




The most common mistake is treating AI adoption like a software rollout.

Buy the tool. Configure it. Train the team. Go live. Move on.

That works for functional software. A new accounting system. A project management platform. People grumble and eventually comply because they have to.

AI is different. Not because of what it does, because of what it represents to the person being asked to use it.

Some employees will dive in immediately. They'll see right away what it can take off



“—
| ***You can mandate usage.
You cannot mandate
trust.***

their plate, the tedious, the repetitive, the work that was never a good use of their time. These people don't need to be convinced. They become advocates on their own.

Others will be unsettled. And they won't say so out loud. They'll think it: will this replace me? If the machine can do what I do, what exactly am I here for?

Both responses exist in almost every organization. The mistake is treating the first group as the norm and the second as an obstacle.

The skeptics aren't obstacles. They're often your most experienced people - the ones with the deepest institutional knowledge, the ones who have seen enough technology promises to know the difference between a tool and a transformation. They deserve a real conversation, not a talking point about efficiency and innovation.

One of the most effective things I've seen a leader do is formalize that skepticism rather than manage it. Assign your most resistant employees a specific role in the pilot - their job is to find where the AI gets it wrong. To challenge the output. To surface the failure modes the enthusiasts will miss.

The person most motivated to prove AI doesn't work is exactly the right person to make sure it does. And they'll surface problems that would have appeared later, and cost significantly more, without them.

THE SHADOW AI PROBLEM



While organizations are debating whether and how to formally adopt AI, their employees have already decided.

ChatGPT. Claude. Gemini. Copilot. Discovered on their own, in use right now, without oversight or governance or anyone knowing what data is leaving the building.

This is shadow AI. And it is already happening in your organization - whether you know it or not.

The employees doing this aren't being reckless. They found something that works and they're using it. That instinct is right. The absence of any framework around it is the problem.

Customer data pasted into a public interface. Confidential financial information used to generate a report. Proprietary processes uploaded to a tool nobody vetted. All of it happening below the surface, invisible to the people responsible for what happens when it goes wrong.

Banning it doesn't work. It just drives the behavior underground and removes any possibility of managing it responsibly.

“
The right question isn't how to stop shadow AI. It's how to make the shadow unnecessary.
”

A word about the productivity gains - because they are real. The employee using ChatGPT to draft client summaries is genuinely better off than one who isn't. The project manager using AI to accelerate an RFP is saving real time. Dismissing that would be wrong.

The issue isn't whether the tools work at the individual level. They often do. The challenge is that individual productivity gain and organizational data exposure can coexist without anyone realizing it until something surfaces. Managing that risk isn't about slowing your people down. It's about making sure the value they're creating doesn't come with undisclosed liability attached.

WHAT THE RIGHT APPROACH LOOKS LIKE

The organizations getting the human side right share one characteristic above all others: they lead from the front.

Leadership is visibly curious, visibly involved, clear about what they don't yet know. AI isn't something they announced and handed to a technology team. It is something they are navigating alongside their people.

They start with small wins rather than organization-wide rollouts. Find willing participants. Give them a focused use case. Let them succeed and let that success be visible. Nothing moves a skeptic faster than watching a colleague demonstrate that a tool actually works.

They invest specifically in the skeptics. The employees most resistant to AI are often the ones with the deepest knowledge of how the business actually operates. Losing them is expensive. Bringing them along takes patience and genuine engagement. It is worth both.

Technology is the easy part. It can be purchased and deployed on a schedule.



“
Culture cannot be purchased. It has to be led.



WHAT GETTING IT RIGHT **ACTUALLY** LOOKS LIKE



Enough about how AI adoption fails. Here is what success actually looks like in businesses that are willing to do the work.

The pattern is consistent. It is not complicated. But it requires a discipline that is genuinely hard: doing the unglamorous work before the exciting work.

START WITH ONE PROBLEM WORTH SOLVING

The businesses making real progress don't start with a platform decision. They start with a question: what specific problem are we actually trying to solve?

Not "we want to be more efficient." Not "we want to do something with AI." Those aren't problems. They're aspirations - and AI can't solve an aspiration.

A real starting point looks like this: a process your team performs dozens of times a week that shouldn't require a human every time. A decision that takes days when it should take hours because the right information is never in one place. Work that is repetitive, predictable, and consuming time from people who should be doing something more valuable.

The specificity matters. AI isn't a general solution. It's a set of capabilities - some powerful, some overhyped - and which ones matter entirely depends on the problem you're actually trying to solve. Getting that match right is the work.

Worth asking alongside "what problem are we solving?" is "what does this replace?" The organizations getting this right aren't just adding capability - they're retiring the

manual process or workaround that AI is making obsolete. If AI ends up running on top of all the same broken workflows, the complexity doesn't go away. It just gets a more sophisticated interface.

The most common mistake at this stage is scope creep. Someone identifies a problem, gets excited, and the vision expands before anything has been proven. Six months later there's a complex implementation that hasn't delivered a single measurable result.

“
Start narrow. Prove it works. Then expand.
”



DEFINE WHAT SUCCESS LOOKS LIKE BEFORE YOU START

This sounds obvious. It is surprisingly rare.

Before anything launches, there needs to be a clear answer to one question: how will we know if this worked?

Not a vague improvement in efficiency. Something measurable. Hours saved per week on a defined task. A reduction in errors on a specific process. Work recovered and redirected somewhere more valuable.

Without a baseline there is nothing to evaluate. And without something to evaluate, the next decision has no foundation.

Measure before the tool goes live - not after. Measuring against a memory of how things used to work produces numbers nobody trusts.

The discipline is straightforward: define success before you start, and measure against it when you finish.

RUN A FOCUSED PILOT

The 30-day pilot is the most underused tool in AI adoption.

A good pilot is simple: one AI capability, one well-defined process, a small group of willing participants, a defined endpoint. Not a proof of concept for the board. Not a vendor showcase. A real test in real conditions with real data.

A few things matter in the setup. The process should have clean, consistent inputs. If it doesn't, you'll discover your data foundation problems the hard way. Include at least one skeptic alongside the early adopters. The skeptic's questions will surface the problems the enthusiasts overlook.

Most pilots like this are operational in under thirty days. Organizations that spend six months planning a pilot before running it aren't being careful.

They're delaying the moment of truth.

“
Run the pilot. Get the data. Make the decision.
”

The question most executives ask at this point is reasonable: how long does this take, and what does it cost?

The real answer depends on what we



find. What I can say is that the foundation assessment and the first pilot don't run sequentially. They run alongside each other. You don't fix everything and then start. You start, and the work of closing the right gaps becomes clear as you go. Most organizations are seeing initial pilot results within sixty to ninety days of the first conversation.

MEASURE, DECIDE, AND EXPAND



When the pilot ends, two outcomes are possible. Both are genuinely useful.

It worked. Results met or exceeded the baseline. The decision is straightforward: expand, build on what was learned, move to the next use case with more confidence behind you.

It didn't work. Results fell short, the process was more complex than anticipated, or the data foundation wasn't ready. That is not a failure. It's a diagnostic - direct, specific information about exactly what needs to be resolved before AI can deliver value here.

The businesses that treat a failed pilot as a failure miss the point. The ones that treat it as a diagnostic come out of it better positioned than when they started.

What doesn't work is the pilot that never ends. The evaluation extended indefinitely because nobody wants to make a call. The implementation that stays in "testing" for a year because the results were ambiguous and the organization can't commit or move on.

In either direction, a decision is better than an indefinite hold.

THE COMPOUNDING ADVANTAGE

There is a dimension to AI adoption that almost never appears in vendor materials. I've started to think it's the most important one.

Once an organization has run one successful pilot, something changes.

The organization that runs one successful pilot doesn't just have one working application. It has a team that's been through the process and can run the next pilot faster, with more confidence and less friction. Leadership has seen real results and is willing to invest further. And the organization has something no external case study can provide - internal proof that AI actually works here, in this business, with these specific people.

What compounds isn't just capability. It's organizational intelligence. Every pilot builds pattern recognition - about which processes respond well to AI, which don't, where the data gaps are, and how to close them. That knowledge doesn't exist in a vendor's documentation. It gets built internally, one iteration at a time, and it becomes genuinely difficult for later movers to replicate.

The organizations that wait aren't standing still. They're falling further behind a moving target. And unlike most competitive gaps - which close when the laggard eventually invests - an AI maturity gap widens with time because the leaders keep compounding while the followers are still trying to clear the starting



conditions. The businesses starting now aren't just getting ahead of the ones that wait. They're building a capability that accelerates over time. The gap between them and the organizations still waiting grows every quarter.

“
*The perfect moment
is not coming.*

THE QUESTIONS WORTH ASKING BEFORE YOU START

Every business considering AI reaches the same moment: someone in the room asks whether it's time to move forward.

What happens next determines almost everything.

In most organizations that question gets answered with a vendor evaluation. Platforms get demoed. Features get compared. A purchase decision gets made. And the questions that actually matter - the ones about readiness, about foundation, about people - never get asked.

Before working with any organization on AI, I ask a different set of questions. Not because the answers need to be perfect - they rarely are. But because discovering the gaps before you start is far less expensive than discovering them after.

Here is what that conversation looks like.

In practice, most leadership teams can work through these in about twenty minutes. The answers are rarely surprising. What's surprising is how rarely anyone has asked them before.



ON YOUR DATA

Can your leadership team agree on a **single source** of truth for the information that drives your most important decisions?

This is the first question I ask - and the answer tells me most of what I need to know. If different teams maintain different numbers, if the CRM and the ERP tell different stories, if people keep their own spreadsheets because they don't trust the system - that gap doesn't disappear when AI arrives. It gets amplified.

Do your **systems actually talk** to each other or does information move between them through manual steps, workarounds, and people who figured out a better way years ago?

AI needs connected data. Give it fragments and it produces fragmented output.

How long does it take to get a **reliable answer** to a basic data question?

If the answer comes in days, or arrives with caveats about which version of the numbers was used - the foundation work isn't done.



ON YOUR PROCESSES

Could someone new to your organization follow your **most important processes** without needing to ask someone who's been there long enough to know?

If the real process lives entirely in people's heads rather than anywhere it can be read - AI can't access it. It can only work with what's documented.

Where is your organization **losing the most time** to work that is repetitive, predictable, and shouldn't require a senior person every time?

These are the processes AI is best suited for. If your team can't identify them quickly, that's worth understanding before any tool evaluation begins.



And where would better information, delivered faster, **change the quality** of a decision your leadership team makes regularly?

This is where AI delivers its highest value - not task automation but decision improvement. Organizations that miss this question default to efficiency plays and leave the more valuable opportunity untouched.

ON YOUR PEOPLE

Does your leadership team have a **real point of view** on AI - not the vendor talking points, but an actual position on what it means for your specific business?


If not, that's not a reason to delay. It's a reason to have a different conversation first - one about what AI actually is and what it would mean here, before anything gets purchased.

Do your employees know AI is coming - and do they **understand why**?

The organizations that handle this well treat it as a leadership communication, not an HR announcement. The message comes from the top. It's clear about what will change. It creates space for the questions people are afraid to ask out loud.

Is there a group of **early adopters** in your organization - people curious enough to experiment and credible enough with their peers that their experience will pull others forward?

If not, finding and supporting that group is one of the highest-value investments a leadership team can make before a formal initiative begins.



ON YOUR EXPECTATIONS

Are you **expecting AI to solve a problem** that is actually a process problem, a data problem, or a people problem in disguise?

This is the most important question on this list. AI doesn't fix underlying problems. It amplifies what's there in both directions. The organizations I've watched fail with AI almost always skipped this question. The ones succeeding almost always didn't.

Are you **prepared to act** on the pilot result - whatever it turns out to be?

A pilot that confirms what you expected is useful. One that challenges it is more useful. The organizations getting the most value from AI treat every result as information worth acting on.

And finally: is the decision to move forward being **driven by clarity** about what you're trying to achieve - **or by anxiety** about being left behind?

Both produce action. Only one consistently produces results.

If the real answer to most of these questions is "I'm not sure" or "probably not" - that isn't a reason to wait. It's the exact starting point that



makes a first conversation useful.

The questions aren't a prerequisite. They're a diagnostic. Knowing where the gaps are before anything is decided or purchased is the whole point. And that's a conversation we can have before you've committed to anything.

If any of these questions gave you pause, that's worth thirty minutes.

A NOTE ON VENDOR NEUTRALITY

If you've read this far and noticed that it hasn't recommended a single AI platform or vendor - that's not an oversight. It's the whole point.

The AI market is moving faster than any organization can track. New platforms emerge monthly. What seemed genuinely differentiated six months ago is now table stakes. The vendor that looked dominant last year is being challenged by tools that didn't exist when this paper was first drafted.

In that environment, a technology partner with a stake in a particular outcome is not a partner. It's a salesperson with a longer introduction.

I am not describing bad actors. I am describing structural incentives. When a firm's revenue depends on driving volume toward specific platforms, when its credibility is tied to certain certifications, the advice it gives is shaped by those dependencies - whether the individuals giving it intend that or not. The incentives are built into the model.

Vendor neutrality is the alternative. In the current AI market - where most products are built on the same handful of foundation models regardless of their branding — it may be the most valuable thing a technology partner can offer.

What does it actually look like?

The recommendation starts with the client's situation, not the partner's certifications. It means being willing to say "this tool is not right for you" even



when the opposite answer would be more profitable. Success is measured by whether the client's business improved - not by whether a platform was successfully deployed.

It also means saying something most vendors never will: you might not be ready yet. The foundation work needs to happen first. The timeline is too aggressive. The problem you want to solve with AI is actually a process problem that no tool is going to fix.

A vendor with a quota is structurally unlikely to have that conversation. A genuine partner - someone whose success is defined entirely by the client's outcome - should always be willing to have it.

Ask any executive who has been through a failed implementation. They can tell you which one they had.

A FINAL THOUGHT

The AI Readiness Gap is real. It is widening. The businesses that close it now will build an advantage that compounds in ways later movers will find genuinely difficult to reverse.

Closing it doesn't require perfection. Your data doesn't need to be immaculate before a pilot starts. Your processes don't need to be fully documented. Your team doesn't need to be uniformly enthusiastic about change.

What it actually requires is a clear-eyed look at where things stand. Where your foundation is solid and where it isn't. Which employees are ready and which ones need more time and a different kind of engagement. What problem is worth solving and whether your organization is genuinely prepared to do what solving it requires.

And the right partner. One who asks the hard questions before you've committed to anything. Someone who measures success the same way you do — not by whether a deal got signed, but by whether AI actually worked.

The gap is closeable. But it won't close on its own.

Five years from now, the AI conversation will be largely over. The separation between organizations will have already happened - not between the ones that adopted AI and the ones that didn't, but between the ones that built the



operational foundation to absorb it and the ones that automated their existing dysfunction.

That separation is happening now. Which side of it your organization ends up on is still a choice.

“ —
The question is whether you're willing to make it deliberately.

About the Author



Jeremy Kushner is the CEO of BACS Consulting Group, a managed IT and technology consulting firm headquartered in Santa Clara, California. For more than fifteen years, BACS has helped growing businesses in life sciences, financial services, construction, and many other industries realize better outcomes through technology. BACS works with companies across the country as a vendor neutral technology partner - focused exclusively on what works for the client's business, not what works for a particular platform or product.

If what you've read here resonates - and especially if the questions from earlier left you with more questions than answers - the most useful next step is a conversation, not a proposal.

I have these conversations with executives regularly. No deck or agenda. Most leaders walk away with a clearer sense of where their real obstacle actually is - whether it's data, process, people, or simply a question of timing.

That clarity is worth thirty minutes.

BACSI.COM | INFO@BACSI.COM | 650.887.4601