withum

Five Foundations Readiness Checklist

A 10-minute self-assessment for AI readiness

Instructions

Answer each question honestly based on your organization's current state - not where you plan to be or where you think you should be.

Scoring: Most questions are "Yes/No." For Foundations 4 and 5, there are a few diagnostic questions before the scored items (marked with checkboxes). Answer those to help interpret the scored questions in your context.

Score 1 point for each "Yes" answer on questions with the Yes/No table.

Foundation 1: Clear Direction

Can your organization answer these questions with specifics?

QUESTION	YES	NO
If your CEO asked you right now, "What specific number will AI improve?" could you answer with a metric and target? (e.g., "DSO from 45 to 35 days" or "close cycle from 10 to 7 days" - not "efficiency" or "better insights").	☐	☐
If AI saved your team 10 hours/week, do you know specifically what they'd do with that time? (Not "strategic work" or "customer relationships", but actual tasks like "analyze pricing variance" or "proactive collections calls to top 20 customers").	☐	☐
If you deployed AI for 90 days, how would you know if it worked or failed? (What metric moves? By how much? Who measures it? Not "user feedback" or "adoption rates"—an operational metric someone reports to the board).	☐	☐
If your staff asked "Should I use AI for [task X]," would you know immediately whether to say yes or no? (Without needing to "think about it" or "discuss with the team"— you have a clear framework for what's priority vs. nice-to-have).	☐	☐

Foundation 1 Score: ____ /4

withum

Foundation 2: Safe Usage

Do employees know the rules, or are they guessing (and hiding their AI use)?

QUESTION	YES	NO
You have written AI usage guidelines (doesn't have to be 47 pages, can be three clear boundaries), and employees actually know where to find them.	☐	☐
Employees know what they can do with AI without asking permission first (internal drafts, research, summarizing documents).	☐	☐
Employees know what requires manager or compliance approval before using AI (client-facing content, financial data, anything that goes external).	☐	☐
Employees know what's completely off-limits (customer personal info in unapproved tools, confidential data in external AI systems, autonomous decisions above a certain threshold).	☐	☐
When AI produces questionable output, employees know who to ask and won't get in trouble for asking.	☐	☐

Foundation 2 Score: ____ /5

Score 0-2?
You're likely seeing either chaos (people using AI in risky ways you don't know about) or paralysis (nobody does anything because they're afraid to ask). Write three boundaries. That's it.

withum

Foundation 3: Training Your Team

Is your team prepared to work effectively with AI?

QUESTION	YES	NO
Employees have been trained on when to trust AI output vs. when to verify.	☐	☐
Different verification standards exist for different work types (routine vs. client-facing vs. strategic).	☐	☐
Employees know how to give feedback when AI produces poor results.	☐	☐
There's training on what to do with time saved (not just how to use tools).	☐	☐

Foundation 3 Score: /4

Foundation 4: Data Foundation

Is your data clean enough to support operations today?

First, what systems do you actually use? (Helps you interpret questions below):

- ☐ Enterprise ERP (SAP, Oracle, NetSuite, Microsoft Dynamics)
- ☐ Mid-market ERP (Acumatica, Epicor, Sage Intacct)
- ☐ Accounting software + other tools (QuickBooks, Xero + spreadsheets/CRM)
- ☐ Mostly spreadsheets, email, and tribal knowledge

withum

QUESTION	YES	NO
CUSTOMER DATA: Customer records are clean: no duplicate entries, contact information is current, and collections reach the right people.	☐	☐
PRODUCT/PRICING DATA: Product categories and pricing live in a single system of record (not manually updated spreadsheets), preventing pricing errors and inventory confusion.	☐	☐
FINANCIAL DATA: All departments use the same numbers for key metrics: no reconciliation needed because Sales and Finance already agree on revenue figures.	☐	☐
DUPLICATE RECORDS: You rarely find duplicate customer, vendor, or product records (less than one per month), and when found they're fixed quickly.	☐	☐
DATA CLEANUP: You have a regular process (not just "when we notice a problem") to identify and fix data quality issues before they cause operational problems.	☐	☐

Foundation 4 Score: ____ / 5

Score 0-2?
These data problems cost money today through higher DSO, pricing errors, and extra close time. Fix them for operational ROI. AI can wait.

Foundation 5: Workflow-Embedded Deployment

Will AI tools fit into how your team actually works?

First, list the 3-5 systems your team uses every day:

1.	(e.g., QuickBooks, Excel, Salesforce, email, etc.)
2.	
3.	
4.	(optional)
5.	(optional)

withum

QUESTION	YES	NO
If you deployed AI today, it would appear INSIDE the systems above (not require opening a separate browser tab or logging into a different portal).	☐	☐
Your team wouldn't need to stop what they're doing, switch to a different application, copy data over, then copy results back.	☐	☐
AI recommendations would show up on the screens your team already looks at every day (not in a separate system they'd have to remember to check).	☐	☐
Before selecting any AI tool, you've identified where your team actually does their work and confirmed the tool integrates there.	☐	☐

Foundation 5 Score: ____ / 4

Calculate Your Total Score

FOUNDATION	YOUR SCORE	MAX SCORE
Clear Direction		**4**
Safe Usage		**5**
Training Your Team		**4**
Data Foundation		**5**
Workflow-Embedded Deployment		**4**
TOTAL		**22**

Questions about your score? A short discussion can help clarify which foundation to address first and where risk may be concentrated. Contact us to talk through your results.

withum

What Your Score Means

0-7 Points: Not Ready Yet (But That's Fixable)

You have gaps across multiple foundations. Deploying AI now would likely repeat the pattern where 95% of AI pilots fail—not because the technology doesn't work, but because the operational foundations aren't there.

What to do: Look at which foundation scored lowest. That's where operational problems are costing you money TODAY (not someday when you deploy AI).

Foundation 4 (Data) lowest?	Customer data chaos costs you through higher DSO, pricing errors cost revenue, conflicting financial numbers extend your close. Fix ONE of these for operational ROI. AI-readiness is the side effect.
Foundation 1 (Direction) lowest?	Pick one metric to improve (DSO? Close cycle? Forecast accuracy?) before buying any tools. Without a target, you can't measure success.
Foundation 2 (Safe Usage) lowest?	Write three boundaries. Not 47 pages. Three rules: what's allowed, what needs approval, what's off-limits.

Start with operational problems that cost money today. AI comes later, after the basics work.

8-14 Points: Foundations Partially Built

You have some pieces in place but gaps remain. This is where most mid-market companies land — not failing, not fully ready, somewhere in between.

What to do: Look at your lowest-scoring foundation. One weak foundation undermines everything else.

Foundation 1 (Clear Direction) lowest?	Define one specific metric to improve before doing anything else. "Reduce DSO from 45 to 35 days in 90 days" beats "experiment with AI" every time.
Foundation 4 (Data) lowest?	You don't need perfect data everywhere. Fix the data needed for your first use case. DSO reduction? Clean customer data. Forecast accuracy? Clean product data. Targeted fixes, not perfection.
Foundation 2 or 3 lowest?	Your team either doesn't know the rules (chaos) or doesn't know when to trust AI output (wasted investment). Both are fixable in a week with clear guidelines.

Pick the weakest foundation. Fix it for operational reasons (it's costing money today). AI-readiness improves as a side effect.

15-18 Points: Ready for Focused Deployment

Your foundations are solid enough to support AI deployment without blowing up. You're in the top 13% of companies in terms of readiness.

What to do: Pick ONE operational metric to improve. Deploy AI to improve that specific metric. Measure results in 90 days.

Examples: DSO from 45 to 35 days (AI-assisted collections prioritization) - Close cycle from 10 to 7 days (automated reconciliations, anomaly flagging) - Forecast accuracy from 70% to 85% (demand forecasting on clean product data)

Start small. Prove ROI. Learn what works. THEN expand to other use cases. Don't try to "transform operations" all at once — enterprises tried that and 95% failed.

19-22 Points: Ready to Scale (Rare)

Exceptional readiness across all foundations. You're positioned to move beyond single-process improvements to more ambitious initiatives.

What to do: You can handle cross-functional AI deployment (order-to-cash automation, continuous close) because your foundations won't collapse under the complexity.

But here's the question: Should you? Or should you get more operational leverage from narrow, boring improvements first?

Sometimes the highest ROI is still the unglamorous stuff: better collections prioritization, automated reconciliations, predictive reorder points. Just because you CAN do the ambitious thing doesn't mean it's the right move.

Consider: What operational problem costs you the most money today? Start there, even if it's not the sexy AI project.

Questions about your score? We're happy to talk through what you're seeing. Sometimes a 15-minute conversation clarifies which foundation to tackle first.

Next Steps

Want the full picture? Explore the full AI Readiness series for detailed analysis of why each foundation matters.

withum

Foundation-Specific Guidance

If Foundation 1 (Clear Direction) is your lowest:

Start here. Without clear direction, teams experiment randomly, value goes unmeasured, and adoption stays low because no one knows what "success" looks like. Define specific operational metrics before buying any tools.

If Foundation 2 (Safe Usage) is your lowest:

You're seeing one of two problems:

Chaos: 57% of employees hide their AI usage because they don't know the rules. Result: compliance risks you don't know about, data leaking into systems you don't control, inconsistent quality.

Paralysis: Nothing happens without committee approval. People are afraid to try anything. Innovation dies.

The fix: Write three boundaries (not 47 pages):

GREEN ZONE	What's allowed without asking (internal drafts, research, summarizing — anything where a human reviews before it goes external)
YELLOW ZONE:	What needs approval (client-facing content, financial data, connecting AI to internal systems)
RED ZONE:	What's off-limits (customer personal info in unapproved tools, confidential data in external AI)

That's it. Employees can remember three boundaries. They can't remember 47 pages.

If Foundation 3 (Training) is your lowest:

Your team either trusts AI blindly (dangerous) or doesn't trust it at all (wasted investment). Both are expensive.

The fix: Train judgment, not prompting. Different work needs different trust levels:

- Routine internal work: AI assists, human spot-checks (doesn't need 100% review)
- Client-facing work: AI drafts, human reviews EVERY output before it goes out (no exceptions)
- Strategic decisions: AI researches options and surfaces information, human decides

Stop teaching "how to write better prompts." Start teaching "when do I trust this output vs. when do I verify it?" That's the skill that matters.

withum

If Foundation 4 (Data) is your lowest:

This is the most common blocker — and the good news is these problems cost you money whether AI exists or not. Don't try to clean up all data. Pick one operational problem that's expensive today:

- **DSO too high?** Fix customer data (duplicates, outdated contacts, wrong credit terms)
- **Pricing errors losing revenue?** Fix product/pricing spreadsheets (get them into a system with one current version)
- **Close takes too long?** Fix financial data fragmentation (pick ONE number everyone uses for key metrics)

Targeted cleanup for one specific problem. Measure the operational ROI. AI deployment becomes possible as a side effect.

If Foundation 5 (Deployment) is your lowest:

If your team must stop what they're doing, open a separate system, copy data over, and copy results back - they won't use it consistently. Adoption dies from friction.

Before selecting ANY AI tool, answer:

- **Where does my team spend 80% of their time?** (QuickBooks? Excel? Salesforce? Email?)
- **Will this AI tool work INSIDE those systems, or require switching to a separate portal?**
- **If it's a separate portal, is the value worth the friction?** (Usually no.)

Deploy AI where work already happens, not where you wish work happened.

Before You Proceed: What Has to Be True

You've scored yourself. Now the honest question: is this approach right for you, or are you optimizing for something else?

This works when you're ready to fix operational problems for operational ROI, with AI-readiness as the side effect. It doesn't work if:

You're chasing AI-readiness as the goal (not operational improvement)

If you need to "show AI progress" to your board but aren't willing to fix the underlying operational problems those foundations revealed, this approach won't work. We fix operations first. AI-readiness happens as a side effect.

What to do instead: Pick the operational problem from your lowest-scoring foundation that's costing you the most money today. Tell your board: "We're fixing [the specific operational problem]. Measured ROI in 90 days. AI deployment becomes possible once the foundation is solid."

Examples by foundation

FOUNDATION 1	FOUNDATION 2	FOUNDATION 4
"We're defining clear success metrics for our three priority initiatives before investing in tools."	"We're establishing AI usage guidelines, so employees know what's allowed vs. what needs approval."	"We're fixing customer data mismatches / pricing spreadsheet chaos / financial reconciliation delays."

That's a story you can defend whether AI works out or not.

You need a prescription without diagnosis

We can't recommend "fix these three things" without examining the actual problems. If you need answers before letting anyone look at how things currently work, we're not the right fit.

What to do instead: Start by measuring the problem your lowest foundation revealed:

FOUNDATION 1	FOUNDATION 2	FOUNDATION 4
Can your leadership team write down, right now, the specific metric AI should improve and by how much? If not, that's the starting point.	Survey 10 employees: "What AI tools are you using, and how do you know if that's allowed?" If you get 10 different answers, you've found your problem.	Export your customer records from CRM and ERP. Count how many match vs. duplicates/mismatches. Calculate match rate. If >95%, data is clean. If <80%, you've measured the problem size.

Once you've measured it, you can decide whether to fix it yourself or get help.

You believe your data is already clean (without having measured it)

If you THINK customer data matches between systems but haven't actually calculated the match rate, you're either in the top 5% of companies (genuinely rare) or you haven't looked closely yet. The Acme Analytics disasters usually start with "our data is pretty clean."

What to do instead: Test your assumption. Export 100 customer records from your CRM. Export the same 100 from your ERP. See how many matches perfectly (name, address, contact, credit terms). If match rate is >95%, you're right—data is clean. If it's <80%, that's your starting point. Measure before you assume.

withum

You need complete results in 30 days for a board meeting

Real operational improvements take 8-12 weeks for focused problems. You'll see progress in 30 days (analysis complete, process designed, pilot started). But if you need the COMPLETE fix delivered before your next board meeting, you're optimizing for theater instead of results.

What to do instead: Tell your board you're running a 90-day operational improvement pilot on your lowest-scoring foundation. Show progress at each stage:

30 days	Analysis complete. "Here's the problem size, here's what it costs us, here's the approach."
60 days	Process built, pilot tested. "Here's the new process, here's pilot results, here's measured impact."
90 days	Full rollout, ROI measured. "Here's operational improvement, here's what we learned about AI-readiness."

That's a defensible timeline regardless of which foundation you're fixing.

You're shopping for the cheapest bid

Customer master data cleanup costs $8K-$15K depending on how messy it is (analysis + reconciliation process + temp worker or consultant time). ROI is typically 15-30x when it's actually fixing a DSO problem. But if your buying criteria is "lowest cost proposal," we're not competitive.

What to do instead: Calculate what the operational problem costs you today. DSO at 45 days instead of 35 days? That's $X in delayed cash flow. Pricing errors from spreadsheet chaos? That's $Y in lost revenue. If the problem costs less than $50K/year, hire a temp for $25/hour to reconcile data manually using internal knowledge. If it costs more than $50K/year, the ROI justifies proper help.