

EXECUTIVE GUIDE · AI ADOPTION SERIES

Are you ready for AI?

Organizational and data readiness, and how to choose a solution that fits the reality you actually have.

THE GUIDE FOR

Proposal, capture and BD leadership

READING TIME

20 minutes

WRITTEN BY

VisibleThread

INCLUDES

The AI Readiness Assessment, a scored 25-statement worksheet

AI adoption rarely stalls because a team picked the wrong tool.

It stalls because the organization, the data, and the process underneath were not ready for it.

This is Part 1 of a two-part series on putting AI to work in bid and proposal organizations. It covers the work before you spend anything: diagnosing what's holding your team back, and using that diagnosis to evaluate a solution.

The order matters. Your readiness diagnosis turns a vague comparison into pointed vendor questions. Part 2 picks up after the purchase decision, with workflow design, pilots, and change management.

Selected sections carry a **Do this next** prompt: a practical action or exercise to apply the idea to your own organization. Not every section has one. Every chapter closes with a short checklist you can run against your own organization.




THE IDEA THAT RUNS THROUGHOUT

The most effective AI workflows are not purely generative. They combine deterministic processes, generative AI, and human judgment into one structure your team can trust and follow.

Part 1 gets you ready for it. Part 2 builds it.

THE ADOPTION LOOP



PART 1 · THIS GUIDE

PART 2 · MAKING IT WORK

Contents

Part 1 covers readiness and selection. Part 2 covers everything after the purchase decision.

PART 1 · YOU ARE HERE

Are you ready?

Organizational and data readiness, and how to evaluate a solution against it. Includes the scored AI Readiness Assessment.

PART 2

Making it work

Designing hybrid workflows, scoping a pilot, training, change management, and measuring what matters.

NO.	CONTENTS	PAGE
01	Organizational and data readiness Why deployments fail, the four pillars, and what an AI-ready library looks like.	4
02	Selecting the right solution A readiness diagnostic, the constraints that decide your shortlist, and seven questions to ask.	12
03	Bringing it together Where to go from here and what to do if you are not ready.	22
A	The AI Readiness Assessment A scored, 25-statement worksheet to run against your own organization.	24

HOW TO USE THIS GUIDE

STEP ONE

Read it once

Twenty minutes, front to back, for the shape of the problem.

STEP TWO

Then score it

Run Appendix A with three to five people and compare.

STEP THREE


Then act on it

Take the Do this next steps and the two checklists into your own meetings.

Organizational & data readiness

Most organizations already have an enormous amount of knowledge. The real challenge is structure, accessibility, and trustworthiness.

This is the step teams are most tempted to skip. Data readiness is not about having more information; it is about whether the information your team depends on is structured, accessible, reliable, and governed enough for automation to use. Without that, automation does not fix disorder. It scales it.



Why AI deployments actually fail

80%

of AI projects fail by some estimates, roughly twice the rate of comparable non-AI IT projects.

ESTIMATES CITED BY RAND

The ratio is the useful part. If AI projects fail at double the rate, the causes have to be specific to AI rather than the usual project risks. Scope creep and budget overruns are not twice as common on AI work.

What is specific to AI is everything around it: the content it reads from, the process it plugs into, and the expectations set before it arrived. Deployments fail when those three things aren't ready, not because the model lacked capability, and all three are within your control to assess before you buy.

Six ways adoption quietly breaks

Before a tool enters the picture, the same patterns show up again and again: six surfaces of the same readiness gaps.

01

Partial adoption

Only a few power users engage, while everyone else keeps the old workflow.

02

Tool fatigue

Too many systems, so teams default back to familiar manual processes.

03

Workaround culture

People export outputs into other tools because the system does not fit their work.

04

Data avoidance

Users steer clear because they know the underlying content is outdated or unreliable.

05

Shelfware

The software is purchased but rarely used after rollout.

06

Demo-only success

The tool shines in a demo but struggles against real proposals.

DO THIS NEXT

Pick the one pattern above that your team is closest to today and write down the single behavior that proves it. That behavior is your symptom. The rest of this chapter is about finding the cause underneath it, because in the 80% that fail, the cause is almost never the tool.

The three root issues

Underneath those symptoms sit three foundational problems. Naming them honestly is the first readiness exercise.

ROOT 1	ROOT 2	ROOT 3 · THE QUIET KILLER
Process immaturity	Data chaos	Misaligned expectations

Of the three, misaligned expectations is the quiet killer. Process and data problems are visible, and a team that has them usually knows it. Expectations are invisible until they are disappointed, and they create the perception that a tool is ineffective when the real issue is organizational readiness.

EXPECTATION	REALITY
Instant productivity	There is an adjustment period
Automatic adoption	Adoption requires change management
AI writes everything	Humans remain in the loop
Immediate ROI	Value grows as processes mature

DO THIS NEXT

Score yourself on process, data, and expectations. The lowest score is where adoption will break first, so that is where to start. Then put these four expectations in front of leadership in writing, because agreement now prevents the “it isn’t working” conversation in month two. Appendix A scores all of this properly.

The four pillars of data readiness

Being “data ready” means the information your team relies on is structured, accessible, reliable, and governed. One definition, how it breaks, and the question to ask internally for each.

01 Content integrity — the quality and reliability of the information feeding your process.

Strong integrity means trusted, validated assets, not archives of old proposals. It breaks with conflicting descriptions and no review cycles — AI reuses whatever it finds, so errors scale fast.

Ask: Is content reviewed on a cycle, approved before entering the library, and version-controlled?

02 Accessibility — how quickly teams and systems can locate and retrieve what they need.

AI needs clear signals to find relevant content. It improves with a consistent taxonomy, centralized repositories, and standard naming.

Ask: Could a new hire find the right past performance example without asking a colleague?

03 Process maturity — whether the work happens the same way twice.

Informal processes give automation nowhere reliable to plug in. Mature processes create clear points where it adds value: intake, drafting, review.

Ask: Could you describe the process to a new hire from a document, not from memory?

04 Governance & ownership — the rules and guardrails that keep AI use responsible.

Who maintains the library, and who approves AI-generated content before it's final? It keeps AI aligned with your standards.

Ask: Is there a defined approval path, and standards for how output is reviewed?

THE OUTCOME

Reliable, AI-ready proposal operations

Weak in one pillar, and automation scales the weakness. Scored in Appendix A, one owner per pillar.


Working backward: map use cases to data

Most teams start in the wrong place, with the tool. Start with the deliverable instead.

A practical readiness exercise is to start from the deliverable you want and work backward to the data and source systems required. Define the output, identify the data it needs, locate the system of record, and confirm the integrity controls around it.

« WORK BACKWARD

START HERE



DELIVERY RUNS FORWARD · CONTENT FLOWS FROM CONTROLS TO OUTPUT »



The same exercise, four more deliverables

Applying the same backward pass to the deliverables that recur most often in a proposal team.



DESIRED OUTPUT	DATA REQUIRED	SOURCE OF TRUTH	INTEGRITY CONTROLS
Past performance example	Metrics, outcomes, customer details	CRM / validated PP database	Evidence validation, customer approval, update cycle
Compliance matrix	Requirement statements, section mapping	Solicitation documents	Deterministic extraction, requirement tagging
Executive summary	Win themes, differentiators, customer priorities	Capture plans, strategy repository	Leadership-approved messaging, controlled access
Amendment updates	Changed requirements, impacted sections	RFP amendments, version-controlled docs	Change tracking, audit trail

DO THIS NEXT

Take your single most common deliverable and fill in one row of this table for it. If you cannot name the source of truth, that is your first content project.

Document-centric vs. AI-ready libraries

The single biggest data-readiness shift is moving from a folder of finished documents to a structured, modular knowledge library. The difference shows up directly in the quality of AI output.

Traditional library document-centric	AI-ready library structured content
STRUCTURE Folders of past proposals, large narrative docs, reuse by copy/paste. CHALLENGE Proof points buried in long documents; hard to locate capabilities; content goes stale silently.	STRUCTURE Modular content blocks, defined taxonomy, rich metadata, evidence separated from narrative. RESULT Precise retrieval; faster assembly; traceable, validated proof points; reusable across opportunities.
AI RESULT Generic or hallucinated content	AI RESULT Grounded, accurate proposal content

Automation scales whatever you give it. Structure the input, and the output follows.

DO THIS NEXT

Open your proposal library and find one reusable section buried inside a finished proposal. Pulling it out into a tagged, standalone block is the first concrete step toward an AI-ready library.

Right-size the framework

The four pillars scale down as readily as they scale up. What changes is who wears which hat.

SMALL TEAM

Assign responsibilities, not committees. One source repository, one workflow, one owner. Roles are hats, not headcount.

MID-SIZE

Shared standards and a repeatable pilot process across proposal teams, with named owners per pillar.

ENTERPRISE

Add business-unit governance, platform controls, common taxonomies, security boundaries, and centralized measurement.

CHAPTER 01 CHECKLIST

Are you ready?

- ☐ We can name the adoption-failure pattern we are most at risk of.
- ☐ Each of the four data-readiness pillars has a named owner.
- ☐ We've scored ourselves against Appendix A, with several people scoring independently.
- ☐ We know our lowest pillar and treat it as the constraint, not average it away.
- ☐ We've started moving to modular, tagged content and mapped a deliverable to its source of truth.
- ☐ Leadership has agreed, in writing, on what AI will and will not do.

Selecting the right solution

Identifying gaps early lets you align internal readiness with vendor capabilities before committing to a solution.

That readiness diagnosis is your biggest advantage in an evaluation, because it turns vague feature comparisons into pointed questions. The four pillars become your scorecard: a vendor either helps with content integrity, accessibility, process maturity, and governance, or it does not.

The presale readiness diagnostic

Four question sets, one goal: do you know your own environment well enough to evaluate a vendor?

Data

Where does proposal content live?
Is reusable content tagged and structured?
Can we identify a source of truth?

WHY IT MATTERS

Content is often technically available but not organized in a way automation can interpret reliably.

Workflow

How does opportunity intake happen?
Who drafts initial content, and who reviews it?
Where do handoffs or bottlenecks occur?

WHY IT MATTERS

These answers reveal where automation would deliver the most value.

Technology

What security or authorization requirements apply?
What must it integrate with: SharePoint, CRM, DMS?
What data residency or compliance constraints apply?

WHY IT MATTERS

In a regulated environment these constraints do more than shape the shortlist. They decide it.

Culture

Are SMEs willing to contribute structured knowledge?
Is leadership aligned on why AI is being considered?
Who will own adoption?

WHY IT MATTERS

Organizations that address cultural factors early are far more likely to achieve strong adoption.

How buying decisions actually get made

Before you evaluate a single vendor, it is worth being honest about how this decision tends to get made, because the pull toward the familiar is stronger than most teams admit.



25%

prefer established vendors
adding AI capability

5%

prefer AI-native entrants, a five-
to-one margin

61%

evaluate both equally, based on
the specific business need

INFUSE, mid-year survey of 310 B2B technology buyers. Brand familiarity was rated very or extremely influential by 45% of buyers.

Familiarity

- Already through security review
- Already integrated
- Already understood by your team

≠ **Fit**

- Built for the work you do
- Answers the question you are asking
- Weighed against the specific business need

Three things that separate the two

That instinct toward the familiar is not irrational: an incumbent is already through security review, already integrated, and already understood by your team. But familiarity answers a different question than fit. The more encouraging finding is that 61% evaluate both equally against the specific business need. So the question to bring to an incumbent is not whether you trust them — it's whether the AI capability is a deliberately engineered part of the platform or a feature added to keep pace with the market.



01

Where it came from

Built, acquired, or wrapped around a general-purpose model? When did it ship, and how many customers are in production rather than in pilot?

02

How deep it reaches

Does it read your existing content, permissions, metadata, and version history natively, or operate on a separate index you must maintain?

03

What it was built for

A capability trained on general business content behaves differently against a 200 page solicitation. Ask which it is, and ask to see the difference.

Ask an AI-native vendor the mirror image of those questions: not whether the AI is good, but whether the surrounding platform, security posture, and support model can carry a regulated program.

DO THIS NEXT

Write down which way your organization is currently leaning, incumbent or AI-native, and why. If the reason is familiarity rather than fit against your four diagnostics, you have found the bias you need to control for.

The constraints that decide the shortlist

Hard constraints that determine which vendors are viable at all, worth resolving before you invest weeks in an evaluation.



01 ○ What is the content?

Sensitivity is set by the most sensitive element in the library: CUI, personal data, export-controlled material.

02 ○ Where can it be processed?

FedRAMP level, DoD impact level, StateRAMP, or your own accreditation. The answer eliminates most of the market.

03 ○ What flows down to you?

Your obligations become your vendor's. CMMC and NIST SP 800-171 put their tools inside your assessed environment.

04 ○ Who else touches it?

Training, retention, subprocessors, personnel, isolation. The terms where AI procurement differs from software.

THE OUTCOME

Viable shortlist

Each gate must clear before the next one matters. The detailed guidance for all five follows over the next pages.

The five gates before you evaluate fit

Each must clear before the next matters — content, hosting, flowdown, model terms, and audit-readiness.

01 What's in your content

Technical volumes can hold CUI, customer data, or export-controlled detail, often unmarked, so the library inherits its most sensitive element. Ask: what's the most sensitive category entering this system?

02 Where it can be processed

The bar may be FedRAMP, a DoD impact level, StateRAMP, or your own accreditation — ask early. Watch for two traps: "in process" isn't authorized, and platform authorization doesn't cover a bolted-on AI feature.

03 What flows down to you

Under CMMC and NIST SP 800-171, tools touching CUI sit inside your assessed environment. Ask what framework a vendor is assessed against, and whether they can produce audit artifacts.

04 Model and data handling terms

Get in writing: is content used to train any model, how long is it retained, which subprocessors sit behind it, and is it isolated from other customers?

05 Records, audit, defensibility

A protest can require you to reconstruct what was submitted and how. AI adds a question: what was machine-generated, from what source, and who approved it?

Start security review before evaluation — it's the most underestimated item in an AI procurement timeline. A vendor who can't clear the gate was never a real option.

Judge the fluency, not just the answers — a vendor who works routinely here answers immediately, with specifics, not adjectives.

DO THIS NEXT

Write down the single hardest constraint you operate under, and lead with it in your next vendor conversation. It will shorten your shortlist in one meeting.

Evaluating vendor claims

Vendor demos run on clean, curated data. Your environment does not. Take each claim, ask implementation questions, and evaluate fit for your specific environment.

THE CLAIM

THE QUESTIONS THAT CUT THROUGH IT

"Our AI drafts full responses instantly"

What data sources does it use, ours or external models? How does it decide what's relevant, and how are conflicting sources handled?

"We integrate with SharePoint"

Read-only or bidirectional? Does it index automatically, recognize metadata, and let us control which libraries are included?

"Implementation is quick and easy"

What assumptions are baked into that timeline? What internal prep, data cleanup, and training is really required?

"Our AI understands requirements automatically"

Which solicitation formats are supported? How accurate is requirement ID, and can we customize classification rules?

A pattern worth noticing

Every one of those questions is answerable in a sentence by a vendor who has done this work, and answerable only in adjectives by one who has not.



A SPECIFIC ANSWER

Names the mechanism, the format, the limit, or the document. You can write it down and hold them to it.

A VAGUE ANSWER

Restates the claim in different words, or moves to a different feature. That is a finding, not a dead end.

A DEFERRED ANSWER

"We can look into that."
Reasonable once. A pattern of it tells you where the product is thin.

DO THIS NEXT

In your next demo, bring one of these claims and refuse to move on until you get a specific, non-marketing answer. Vague answers are data too.

Seven questions to ask every vendor

Beyond feature lists, these questions reveal whether a solution will work in your environment. Ask all seven, in order.

☐ **1** ASSUMPTIONS
What assumptions does this feature make about our data?
Does it need tagged content? Can it interpret unstructured past proposals? How much cleanup is typical?

☐ **2** GROUNDING
Where does the AI get its answers from?
Internal content only, or external models? Can we control which documents are used? Can it cite its source?

☐ **3** WORKFLOW
What workflow does this actually support?
Which lifecycle stage: RFP analysis, drafting, compliance, review? How does it fit existing workflows?

☐ **4** IMPLEMENTATION
What does implementation really require?
Internal prep, integrations, training, and the effort needed from your team during rollout.

☐ **5** TRANSPARENCY
How transparent are the outputs?
Source citations? Traceability back to original content? How easy is it to verify or edit?

☐ **6** CONTROL
What level of control do users have?
Customizable prompts and templates? Guardrails to enforce standards? Permission and admin management?

☐ **7** ADOPTION
What does successful adoption look like?
What results do successful customers achieve, and what does the support model include?

Take this page into your next vendor meeting.

Vague answers are data too.

Demo environment vs. your real environment

Every demo works. The question is whether the tool works where you work.



THE DEMO ENVIRONMENT

Controlled

- Clean curated content
- Structured knowledge base
- Ideal prompts
- Short example documents

>



YOUR REAL ENVIRONMENT

Uncontrolled

- Mixed-quality historical proposals
- Scattered archives
- Inconsistent requirement language
- 200+ page solicitations

IN THE DEMO

Everything in its place.

IN YOUR LIBRARY

Spread across four
systems, none
authoritative.

Evaluate the tool against your reality, not the demo's.

Choosing well

Run these checks before you sign.

- ☐ **We have run** the data, workflow, technology, and culture diagnostics.
- ☐ **We named the vendor** we were predisposed toward, and tested that preference against the diagnostics.
- ☐ **We confirmed the AI capability** sits inside the authorization boundary, not just the platform around it.
- ☐ **Training, retention, subprocessor,** and personnel terms are in the contract, not the sales conversation.
- ☐ **Security review started** before evaluation, not alongside it.
- ☐ **We saw the tool run** against our own documents and pushed every major claim with implementation questions.
- ☐ **We know the most sensitive category** of content entering the system, and whether the vendor can lawfully hold it.
- ☐ **We chose for fit** with our reality and can state our readiness pattern in a single sentence.

DO THIS NEXT

Ask the vendor to run their demo on one of your real, messy documents, not their sample file. How they react tells you as much as the result does.

Almost everything determining your outcome is already within your control

Get the data and the organization ready. Choose a solution that fits your reality rather than the demo.

Notice how little of either is about the tool, which is the same reason the large majority of AI projects fail, by most estimates, while the models keep getting better.

AI does not replace the discipline of good proposal work. It rewards the organizations that already have it, and it exposes the ones that do not. Most of the work in this part can start this week without a purchase order.

Where to go from here

Three things, roughly a half day in total, none of which require a vendor or a budget.

1

Run Appendix A as a working session

Three to five people, scored independently, then compared. You finish with a baseline, a lowest pillar, and a specific first project. If you do only one thing from this guide, do this one.

2

Write your hardest constraint down

An authorization level, an export control restriction, a flowdown requirement. Lead with it in your next vendor conversation and it will shorten your shortlist in one meeting.

3

Put the four expectations to leadership, in writing

The cheapest work in the guide, the most frequently skipped, and the one that prevents the "it isn't working" conversation in month two.

DO THIS NEXT

Those three steps put you ahead of most organizations that have already bought something.

When you are ready to talk to a vendor

Talk to us, but use this guide when you do.

Chapter 2 asks you to push vendor claims until you get specific answers, and to establish where the AI gets its answers. Those questions were not written for other vendors. Apply them to VisibleThread.

A useful first conversation: bring a real, badly formatted solicitation, your Appendix A result (particularly your lowest pillar), and your hardest constraint.

We will show you the tool against your document, not ours. If your score says foundation work comes first, we will tell you that.

[Book a working session »](#)

And if you are not ready

That is a legitimate outcome, and more valuable than a purchase made too early. If Appendix A puts you in foundation-first territory, the highest-return work has nothing to do with AI: structure the library, name the owners, document the process. That work pays for itself whether or not you ever deploy a tool.

The organizations that succeed with AI are not the ones that started earliest. They are the ones that were ready when they started.

The AI Readiness Assessment

A practical baseline for deciding what to fix, what to prioritize, and when to move.



20**MINUTES**

Complete the assessment
against evidence, not intent.

4**PILLARS**

Content integrity · Accessibility
· Process maturity ·
Governance

1**OUTCOME**

A decision about what to do
next and in what order.

The AI Readiness Assessment

Chapter 1 asks you three times to score yourself honestly. This is the instrument for doing it.

Four pillars, five statements each, plus a gate on page 30 that works differently. It takes about twenty minutes. The output is not a grade. It is a decision about what to do next and in what order.

HOW TO RUN IT · SCORE 1 TO 5

**1**

Not true. We have not started.

**2**

Isolated cases. One team or one person does this.

**3**

Inconsistent. It depends who is running the pursuit.

**4**

Mostly consistent. Documented and usually followed.

**5**

Consistently true. Documented, followed, owned, checked.

Score against evidence, not intent. If a practice exists in a document nobody follows, it scores 2. The most common failure of self-assessment is scoring the policy instead of the behavior.

Have three to five people score independently, then compare. Include a proposal manager, an SME, and whoever maintains the content library. Where scores diverge by two or more points, discuss the spread before you average it.

Write the date on it. This is your measurement baseline for Part 2.

DATE _____

01 Content integrity

Is the information feeding your process trustworthy?

1 2 3 4 5

- 1.1

Reusable content is reviewed and updated on a defined cycle, not when someone notices it is wrong.
- 1.2

Reusable sections are stored as standalone, approved blocks, separate from the finished proposals they came from.
- 1.3

Content is approved by a named person before it enters the library.
- 1.4

Every claim and proof point in the library can be traced to a verifiable source.
- 1.5

Anyone on the team can identify the current, authoritative version of a piece of content without asking.

Pillar 1 subtotal (out of 25)

PILLAR 1, REGULATED ENVIRONMENTS · ONE ADDITIONAL CHECK

Score separately and do not add it to the subtotal, because it is a constraint rather than a gradient: content carrying CUI, export-controlled, or customer-restricted material is identified and marked as such. A score of 3 or below limits your options in Chapter 2, because an unmarked library cannot be selectively protected.

/ 5

02 Accessibility *Can people and systems find what they need?*

1 2 3 4 5

- | | | |
|-------|--|--|
| 2.1 | Proposal content lives in a defined set of repositories, not scattered across personal drives and email. | ☐ ☐ ☐ ☐ ☐ |
| <hr/> | | |
| 2.2 | A consistent taxonomy or tagging scheme is applied to reusable content. | ☐ ☐ ☐ ☐ ☐ |
| <hr/> | | |
| 2.3 | Metadata is captured consistently: capability area, customer, date, owner, approval status. | ☐ ☐ ☐ ☐ ☐ |
| <hr/> | | |
| 2.4 | Naming conventions are standardized and actually used. | ☐ ☐ ☐ ☐ ☐ |
| <hr/> | | |
| 2.5 | A new team member could locate the right past performance example without asking a colleague. | ☐ ☐ ☐ ☐ ☐ |

Pillar 2 subtotal (out of 25)

03 Process maturity *Does the work happen the same way twice?*

1 2 3 4 5

3.1 Opportunity intake and kickoff follow a documented, standard process. ☐ ☐ ☐ ☐ ☐

3.2 Requirements analysis and compliance matrix construction follow a repeatable method. ☐ ☐ ☐ ☐ ☐

3.3 Review cycles are scheduled and defined, not assembled per pursuit. ☐ ☐ ☐ ☐ ☐

3.4 Outlines and templates are reused rather than rebuilt from scratch each time. ☐ ☐ ☐ ☐ ☐

3.5 We could describe our proposal process to a new hire from a document rather than from memory. ☐ ☐ ☐ ☐ ☐

Pillar 3 subtotal (out of 25)

04 Governance and ownership

Does someone own it, and are there guardrails?

12345

4.1

The content library has a named owner responsible for its maintenance.

4.2

Each major content area has a named SME accountable for its accuracy.

4.3

There is a defined approval path for content before it reaches a customer.

4.4

Version history is retained and auditable.

4.5

We have standards for how AI output is reviewed, by whom, and against what.

Pillar 4 subtotal (out of 25)

Chapter 1 calls misaligned expectations the quiet killer, and it behaves differently from the four pillars. Pillar scores are gradients: a low score means work to do. This block is a gate. A team can be data-ready and still fail if the organization around it was never prepared for what adoption requires.

	1	2	3	4	5
G.1 Leadership has agreed in writing on what AI will and will not do, including that humans stay in the loop.	☐	☐	☐	☐	☐
G.2 Leadership understands that value grows as processes mature rather than appearing at rollout.	☐	☐	☐	☐	☐
G.3 We have defined what success looks like and how we will measure it.	☐	☐	☐	☐	☐
G.4 There is a named owner for adoption, distinct from whoever owns the software.	☐	☐	☐	☐	☐
G.5 SMEs and proposal staff have been consulted, and are willing to change how they work.	☐	☐	☐	☐	☐

Gate subtotal (out of 25)

IF THE GATE SCORES BELOW 15, STOP

No pillar score is high enough to compensate. An organization that has not agreed on what AI is for, what success means, or who owns the outcome will produce the "it isn't working" conversation in month two regardless of how clean its content library is. Fix the gate first. It is the cheapest work in this guide and the most frequently skipped.

Scoring

PILLAR TOTAL
OUT OF 100

GATE
OUT OF 25

LOWEST PILLAR
WHICH ONE, AND ITS SCORE

TOTAL	BAND	WHAT IT MEANS
80–100	Ready	Deploy broadly. Your constraint is adoption and change management, not data.
56–79	Ready to pilot	Scope a pilot to one team and one deliverable type where your strongest pillar applies. Do not roll out enterprise-wide yet.
32–55	Foundation first	A tool will amplify your current disorder. Spend one to two quarters on your lowest pillar before you buy.
20–31	Not ready	Buying now is how organizations end up in the large majority that fail. Start with content and ownership.

A note on the bands. This is a practical diagnostic, not an externally validated maturity benchmark or certification. The bands are directional: their job is to structure discussion, surface where scorers disagree, and point to the safest next action, not to grade you against an industry standard.

Your lowest pillar governs. If any single pillar scores below 12, treat yourself as one band lower than your total suggests, regardless of how strong the others are.

Averages hide constraints, and these four pillars do not compensate for one another. Immaculate governance over content nobody can find produces the same result as no governance at all. A mature process pointed at an untrustworthy library produces reliable, repeatable, incorrect output. The pillar you score lowest is the one that will determine what your AI deployment actually delivers, which makes it the only sensible place to start.

Reading the shape, not just the number

The pattern across pillars is more diagnostic than the total, and it usually points to a specific first project.

	SHAPE	WHAT IT USUALLY MEANS	START HERE
	High content, low accessibility	You have good material nobody can find.	Taxonomy and metadata. Fastest visible return of any project in this guide.
	High accessibility, low content integrity	You will retrieve the wrong thing quickly and confidently.	Review cycles and approval before anything else.
	High process, low governance	It works because specific people make it work.	Named owners. You are one departure from losing the process.
	High governance, low process	Rules exist that the work does not follow.	Document the actual workflow, then align the rules to it.
	Even and low across all four	No single fix. This is a program, not a project.	One pillar, one quarter, in the order above.

If this is your lowest area: the first 30 days

A score tells you where the problem is. This is what Monday looks like for each one.

WEAKEST AREA	FIRST 30-DAY ACTION
Content integrity	Identify the authoritative source for your top ten reusable sections, archive obsolete versions, and assign each an owner.
Accessibility	Define the repository and metadata scheme for one content type, usually past performance, and tag it.
Process maturity	Document one repeatable proposal workflow end to end, as it actually happens, before automating any of it.
Governance & ownership	Name an AI/process owner, then define the approval and review rules for content entering and leaving the library.
The gate	Establish one use case, a baseline, a success metric, and a written leadership expectation, before any tool conversation.

ABOUT THE AUTHOR

Questions about your own environment? Bring them to Allison.



ALLISON RITZ

Director of Product Marketing,
VisibleThread · CF APMP

A certified APMP member and frequent GovCon speaker, Allison helps some of the largest government contracting teams turn strategic relationships into wins. She wrote this guide from the same conversations she has every week with proposal and capture leaders working through exactly this readiness question.

BI-WEEKLY LIVE SERIES · HOSTED BY ALLISON

VisibleThread Office Hours

Open, no-fluff sessions on winning more government work: AI, better proposals, and smarter pursuits. Come with questions, leave with answers.

Every other Thursday · 30 minutes · live Q&A

Free · replays after

[Save your spot for office hours »](#)