HOW TO AI™

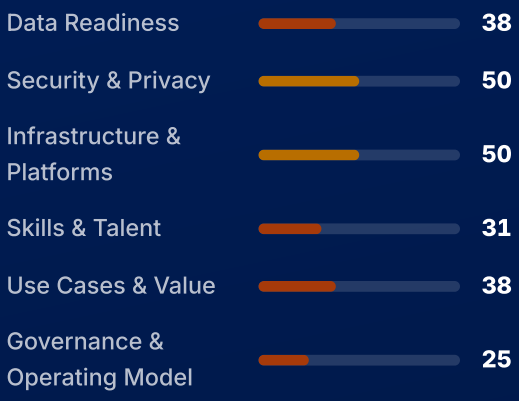
SAMPLE REPORT

AI READINESS REPORT

Example Manufacturing Co.

Your infrastructure is ahead of your guardrails: fix shadow AI and oversharing first, then your service desk goal is very achievable.





Prepared by David Davis · HOW TO // AI
Sample edition

A complete sample of the deliverable
Fictional company; real format

PART —

Contents

01 How to read this report

02 Executive snapshot

03 Executive summary

04 Company profile

05 Scorecard

06 Dimension deep dives — Six chapters, one per dimension

07 Top risks

08 Quick wins

09 The 90-day plan

10 Months 4-12 outlook

11 Budget guidance

12 Measuring success

13 Skills & certification plan

14 A final word

A Appendix: your answers, question by question

About this document. This report was prepared from Example Manufacturing Co.'s intake: 24 scored questions across six dimensions, plus context about the organization, its stack, and its goals. Scores are deterministic — the same answers always produce the same numbers — and the narrative is drafted against the HOW TO // AI readiness framework, then reviewed by David Davis. This sample uses a fictional manufacturer so you can see exactly what you receive.

PART 01

How to read this report

Five minutes here makes the next thirty pages more useful.

The score is deterministic. Each of the 24 questions has five answers describing maturity levels 0-4. A dimension's score is points earned over points possible; the overall score is the plain average of the six dimensions. AI never decides a number in this document.

The rating has a floor rule. If security or governance scores below 40, the overall rating is capped at Emerging regardless of the average — you are not ready if you are leaking data. That rule did not change this rating.

Risks and quick wins come from your answers. Red-flag answers surface first in the risks list, then the lowest-scoring answers. Quick wins are low-scoring questions where the next maturity level is a low-effort fix — weeks, not quarters.

The narrative is where the judgment lives. The analysis, recommendations, plan, and budget guidance are written against your context — your stack, your industry, your regulations, your stated goal — and reviewed by David Davis before delivery.

THE FOUR RATINGS

Not Ready 0-39
Foundations are missing. Adopting AI now adds risk faster than value.
YOUR RATING

Emerging 40-59
Pieces are in place, but gaps will stall pilots or create exposure.

Ready 60-79
Solid foundations. Ready to run and scale targeted use cases.

Leading 80-100
Mature and governed. Focus on scaling value and building.

PART 02

Executive snapshot

The whole report on one page.



Your infrastructure is ahead of your guardrails: fix shadow AI and oversharing first, then your service desk goal is very achievable.

Foundations are missing. Adopting AI now adds risk faster than value.

WORKING IN YOUR FAVOR

- SSO, role-based access, and access reviews are already in place
- Hybrid Microsoft 365 and Azure estate — no new platforms needed
- A concrete, measurable goal: service desk and documentation time
- A Copilot trial in finance is already generating usage lessons

NEEDS ATTENTION FIRST

- Customer and pricing data is going into personal ChatGPT accounts today
- Overshared SharePoint becomes a data leak the day an assistant turns on
- No AI policy, no decision owner, no inventory of AI in use
- No AI skills baseline on the IT team ahead of a build-heavy pilot

THE FIRST MOVES

- Publish the acceptable-use policy that's already in draft
- Give everyone an approved AI assistant, then restrict the rest
- Clean up sharing on your 20 highest-risk SharePoint sites
- Baseline service desk metrics before the pilot starts

DAVID'S TAKE

Your score says Not Ready. Your foundation says otherwise: identity is solid, and you already own the Microsoft tools that fix most of what's broken. The problem is that AI came in the side door. People are pasting pricing into ChatGPT, and SharePoint will hand the wrong file to the wrong person the day you turn Copilot on for everyone. Close those two gaps in the next 30 days, mostly policy and permissions work with no new budget, and the service desk pilot stops being a risk conversation with your CFO and becomes an easy yes.

David Davis

PART 03

Executive summary

Written for you and for the leadership team you'll forward this to.

Example Manufacturing is rated Not Ready at 39/100. That isn't a verdict on your IT team. It reflects that AI is already in use across the company without the policy, visibility, or data controls to make it safe. Your identity controls and hybrid Microsoft footprint are real strengths to build on, and most of what drags the score down is decisions nobody has been asked to make yet, not technology you'd have to buy.

Three issues matter most. Employees are using public AI tools with company data and nothing stops customer or pricing data from leaving; your own intake named this as the thing that worries you, and your answers confirm it's happening today. Your SharePoint and file shares are overshared, which becomes an instant data leak the day an AI assistant is switched on, because assistants surface everything a user can technically open. And nobody owns AI decisions, so there's no inventory, no approved-tool list, and no one who can say yes safely — which is why AI keeps arriving through personal accounts instead of through IT.

The order of operations matters more than the spend. If you fund a service desk pilot first and fix governance later, the pilot inherits every risk on your list: the assistant answers from overshared content, usage grows in tools you don't control, and the first incident sets the program back a year. Run the sequence in this report — policy and cleanup in the first 30 days, foundations in the next 30, pilot in the last 30 — and each step de-risks the next one.

The economics of your stated goal are favorable. Password resets, access requests, SOPs, knowledge base articles, and customer quotes — the processes you called painful — are exactly the workloads where AI assistance is proven and measurable inside a Microsoft 365 and Azure estate. You don't need new platforms; you need the guardrails to use the ones you have, a baseline so improvements are provable, and one pilot scoped tightly enough to show a number your CFO believes.

Your $25K-$100K budget is enough, with room to spare in year one. Nearly everything in the first 30 days is internal time, not spend. The plan in this report holds most of the budget in reserve until the pilot produces evidence, which is also the posture that makes the CFO conversation easy: you're asking them to fund a measured expansion, not an experiment.

PART 04

Company profile

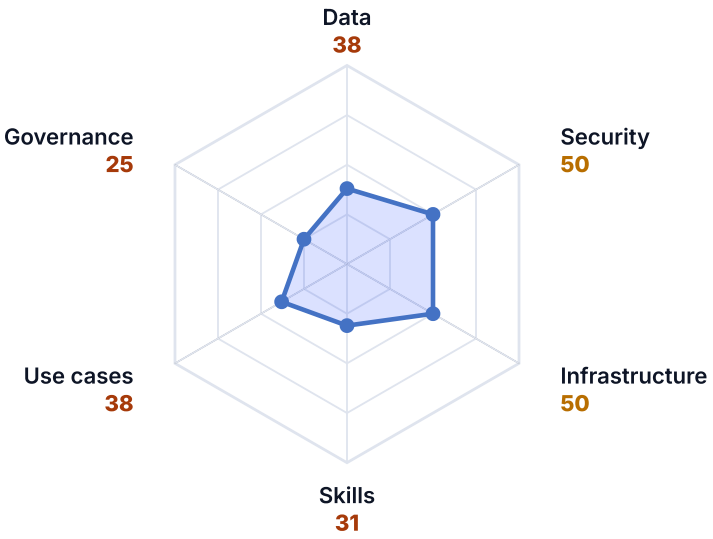
What you told us. Every recommendation in this report is anchored here.

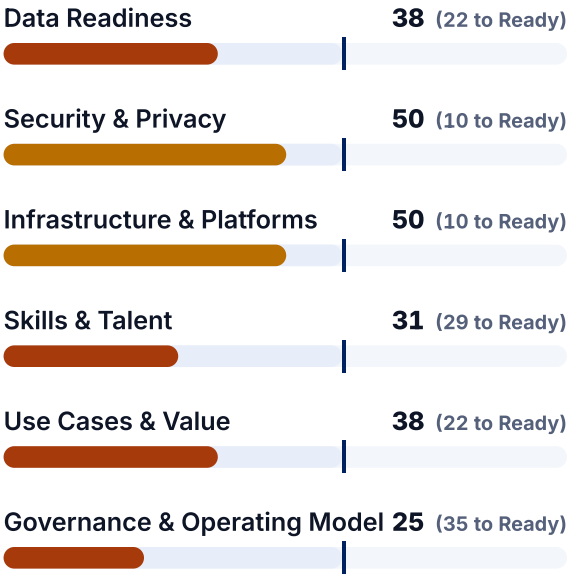
Company name	Example Manufacturing Co.
Number of employees	200-999
Industry	Manufacturing
IT team size	6-20
Primary productivity suite	Microsoft 365
Primary cloud	Azure
Regulations that apply to you	SOX, US state privacy laws
What's the #1 thing you want AI to do for your organization in the next 12 months?	Cut time spent on service desk tickets and internal documentation.
What's your biggest concern about AI adoption?	Employees pasting customer and pricing data into ChatGPT.
Which AI tools do you know people are using today, approved or not?	ChatGPT on personal accounts, Grammarly, a few Copilot trial licenses in finance.
Any AI projects or pilots in flight or planned? Briefly, what are they?	10-seat Copilot trial in finance. Idea stage: a service desk chatbot.
Which 2-3 business processes eat the most time or cause the most pain?	Password resets and access requests; writing SOPs and knowledge base articles; preparing customer quotes.
Rough budget for AI over the next 12 months	$25K-$100K
What would make this report a win for you? (e.g., a board presentation, an approved policy, a pilot plan)	A plan I can take to the CFO to fund a service desk pilot.

PART 05

Scorecard

Six dimensions, scored 0-100 from your answers.





The radar shows the shape of your readiness; the bars carry the exact scores, with the marker at 60 — the Ready threshold. Improving your weakest dimension usually moves the overall rating faster than polishing your strongest one, and if security or governance is below 40, it caps your rating outright.

YOUR ANSWERS AT A GLANCE

Data Readiness	D1 2/4	D2 1/4	D3 1/4	D4 2/4
Security & Privacy	S1 1/4	S2 2/4	S3 3/4	S4 2/4
Infrastructure & Platforms	I1 2/4	I2 2/4	I3 2/4	I4 2/4
Skills & Talent	K1 2/4	K2 1/4	K3 1/4	K4 1/4
Use Cases & Value	U1 2/4	U2 1/4	U3 2/4	U4 1/4
Governance & Operating Model	G1 1/4	G2 1/4	G3 1/4	G4 1/4

Answer level: 0 weakest · 1 · 2 · 3 · 4 strongest

0-39 Not Ready · **40-59** Emerging · **60-79** Ready · **80-100** Leading

Weakest area: **Governance & Operating Model**

38 / 100

Data Readiness

Can AI find, trust, and safely use your data?

Core systems exist, but data ownership is informal and unstructured content is overshared.

Your structured data is in reasonable shape for a company your size: core business apps are in place with some integrations, so the information a service desk assistant or quoting aid would need does exist and can be found. The problems sit one layer down. Nobody owns data quality, so the errors that live in your ERP and CRM today — stale customers, duplicate parts, inconsistent pricing records — will be repeated by any AI that reads them, delivered fluently and with confidence. AI doesn't clean data; it amplifies whatever it finds.

The sharper issue is your unstructured content. You told us permissions on SharePoint and file shares are known to be too broad and haven't been cleaned up. Today that's a dormant problem, because finding an overshared file requires someone to go looking. An AI assistant removes that friction entirely: it reads everything a user can technically access and volunteers it in answers. The pricing sheet buried three folders deep, the HR spreadsheet with a company-wide link, the board deck in a legacy site — all of it becomes one question away. This is the single most important thing to fix before any assistant is connected to company content, and it's why cleanup is in the first 30 days of your roadmap.

Programmatic access is middling: a few one-off APIs and connectors exist, but there's no managed integration layer. That's acceptable for the first pilot — the service desk assistant can start against your ticketing system and knowledge base alone — but each use case after that will re-raise the question, so an integration inventory belongs on the 6-12 month horizon, not in the first 90 days.

STRENGTHS

- Core business apps are in place with some integrations, so structured data is findable.

GAPS

- Nobody owns data quality, so AI answers will inherit today's errors.
- SharePoint and file share permissions are known to be too broad and haven't been cleaned up.

What to do about it

1. Run the sharing and access reports in your Microsoft 365 admin and compliance centers for your 20 most sensitive SharePoint sites, remove 'Everyone' and organization-wide links, and assign a named owner to each site — before any AI assistant is enabled.

2. Name data owners for the ERP, the CRM, and the service desk knowledge base, each with a one-page list of quality rules they're accountable for; these three systems are the ones your pilot will read from.

3. Start sensitivity labeling with the two data types you already know are radioactive — customer records and pricing — rather than boiling the ocean with a full classification program.

4. Before the pilot, have the service desk lead curate the knowledge base: archive superseded articles and mark the trusted ones, so the assistant answers from the good half of the KB and not the graveyard.

5. Defer the integration-layer decision to months 7-12, and capture it as a written architecture question now so it doesn't get decided by accident by the first vendor who shows up.

What we asked

Your answer to each question, scored 0-4. The full maturity ladders are in the appendix.

D1. Where does most of your core business data live today? 2 / 4

"Core apps with some point-to-point integrations; reporting is still manual"

D2. Who is accountable for data quality? 1 / 4

"IT fixes problems ad hoc"

D3. How well organized and permissioned is your unstructured content (SharePoint, Google Drive, file shares, wikis, ticket history)? 1 / 4

"We know there's oversharing but haven't addressed it"

D4. Can AI tools get to your data programmatically? 2 / 4

"Some APIs or connectors in use, built one-off"

50 / 100

Security & Privacy

Can you adopt AI without leaking data or opening new attack paths?

Identity is solid; data protection and AI-specific controls lag behind actual usage.

This dimension splits cleanly in two. Your identity foundation — SSO, role-based access, periodic access reviews — is genuinely strong for a company your size, and it matters more than most teams realize: every AI control you'll deploy, from an approved assistant to DLP policies, hangs off identity. You've already done the hard, slow part.

The exposed flank is data protection in the age of casual AI use. Employees are pasting work data into public AI tools, there's no approved alternative to pull them toward, and your classification program is a policy on paper with labels applied by hand in a few places. That combination means your most sensitive data — the customer and pricing information you specifically told us you worry about — has no technical barrier between it and a personal ChatGPT account. The fix is a one-two punch, and the order is critical: first give people an approved, enterprise-grade assistant so the productivity demand has somewhere safe to go, then restrict unapproved tools through your web filter. Blocking first just moves usage to personal phones where you have no visibility at all.

Vendor risk needs one targeted upgrade rather than an overhaul: your reviews don't yet ask whether a product's AI features retain or train on your data. Every SaaS renewal that crosses your desk this year will have new AI features in it, so this single added question does a surprising amount of work under your SOX and state-privacy obligations.

STRENGTHS

- SSO, role-based access, and periodic access reviews give you a strong base for any AI rollout.
- Some sensitive data already carries labels.

GAPS

- Employees use public AI tools with work data and there's no approved alternative.
- Vendor reviews don't ask whether AI features retain or train on your data.

What to do about it

1. Stand up one approved enterprise AI assistant for every employee inside your Microsoft 365 tenant, with commercial data protection confirmed in writing, and announce it the same week the acceptable-use policy lands.
2. Two weeks after the approved tool ships, move your web filter from monitoring to warn-or-block on unapproved AI sites, and review the logs monthly at the AI council.
3. Extend sensitivity labels and DLP policies in your compliance center to customer and pricing data first, covering email, endpoints, and the approved assistant.
4. Add two questions to your vendor security review: does any AI feature retain our data, and is our data used for training — and re-ask them at every renewal.
5. Fold AI tools into your existing access reviews so an assistant's permissions get recertified on the same cycle as everything else.

What we asked

Your answer to each question, scored 0-4. The full maturity ladders are in the appendix.

S1. Which statement best describes employee use of AI tools today? 1 / 4

"People use public AI tools with work data; we have no controls yet"

S2. How do you classify and protect sensitive data? 2 / 4

"Sensitivity labels applied to some data, mostly by hand"

S3. How mature is your identity and access management? 3 / 4

"SSO, role-based access, and periodic access reviews"

S4. Before adopting an AI tool or AI feature, do you review how the vendor handles your data? 2 / 4

"Standard vendor security review, nothing AI-specific"

PART 06.3 · DIMENSION 3 OF 6

50 / 100

Infrastructure & Platforms

Do you have the platforms to run, integrate, and pay for AI?

A hybrid Microsoft and Azure footprint that can support AI once governance catches up.

Infrastructure is the dimension where you have the least to build. A hybrid Microsoft 365 and Azure estate puts every major AI capability you'll need this year — an enterprise assistant, Azure-hosted models for the service desk pilot, usage dashboards — one configuration decision away rather than one procurement cycle away. Your score here is held back not by missing capability but by missing decisions: enterprise AI is still 'in evaluation,' and in the vacuum, consumer tools have colonized the gap.

The finance Copilot trial is a real asset that isn't being treated as one. Ten seats have been generating usage patterns, prompts, wins, and duds for months. Before you expand or cancel it, mine it: which tasks stuck, which licenses went dormant, what finance users actually asked for. That evidence should drive your assistant rollout sizing, and it's exactly the kind of data-backed decision your CFO will notice.

The one genuinely new muscle you need is cost management for usage-based AI. Your dashboards exist, but nothing is budgeted or alerted for token-metered spend, which behaves nothing like per-seat licensing — it can double in a good month because adoption is going well. Set Azure budgets and alerts before the pilot starts burning tokens, not after the first surprising invoice.

STRENGTHS

- Hybrid SaaS and Azure workloads mean AI services are within reach without new platforms.

GAPS

- Enterprise AI is still in evaluation, so people fill the gap with consumer tools.
- Cost dashboards exist, but no budgets or alerts are ready for usage-based AI spend.

What to do about it

1. Close the evaluation: pick the approved assistant (the obvious candidate is the one already trialed in finance) and set a 30-day decision deadline so the consumer-tool vacuum ends.
2. Write a one-page readout of the finance Copilot trial — usage, wins, dormant seats — and use it to size the wider rollout.
3. Create Azure budgets and alerts scoped to AI resource groups before the pilot's first token is consumed, with a monthly cost review on the AI council agenda.
4. Keep the pilot inside your existing tenancy and region choices so your SOX and privacy posture carries over unchanged, rather than standing up anything new for it.

What we asked

Your answer to each question, scored 0-4. The full maturity ladders are in the appendix.

I1. **Which best describes your infrastructure today?** — 2 / 4

"Hybrid: key SaaS apps plus some cloud workloads"

I2. **What access does your organization have to enterprise AI platforms?** — 2 / 4

"Evaluating or piloting an enterprise AI assistant (e.g., Microsoft 365 Copilot, ChatGPT Enterprise, Gemini, Claude)"

I3. **How automated are your IT operations?** — 2 / 4

"Shared scripts or an automation/RPA tool in limited use"

I4. **Can you see and control what you spend on cloud and SaaS?** — 2 / 4

"Cost dashboards for major services"

PART 06.4 · DIMENSION 4 OF 6

31 / 100

Skills & Talent

Does your team have the skills to adopt, secure, and build with AI?

Curiosity is there, but no training, no AI credentials, and no budget to close the gap.

Your team has informally experimented — most people have tried AI tools on their own — but nothing has been formalized: no credentials, no owner for AI automation skills, no budget, and training approved case by case. That gap has a direct cost on your roadmap: the service desk pilot needs someone who can build and tune it, and today that person doesn't exist on the team. You'd either delay the pilot or outsource its core skill, and renting the one capability you'll need permanently is the expensive choice.

The fix is cheap relative to everything else in this report. A certification sprint — one fundamentals credential for the whole IT team, one deeper engineering path for a single builder — costs exam vouchers and protected study hours, fits inside a sliver of your stated budget, and pays off twice: the team gets a shared vocabulary for every AI decision this year, and you get your pilot builder. It also signals to the team that AI is something they're being invested in, not something happening to them, which matters for the adoption battles ahead.

STRENGTHS

- Most of the team has tried AI tools, so adoption won't start from zero.

GAPS

- No one holds an AI-related certification.
- No one owns building AI automations, which your service desk goal requires.
- Upskilling is approved case by case with no budget.

What to do about it

1. Run a six-week certification sprint: one fundamentals credential for the entire IT team, with two protected study hours per week and exams booked as people hit practice-test readiness.
2. Pick one engineer to own the service desk pilot build and put them on an Azure AI engineering path now, so the skill lands before Days 61-90 of the roadmap.
3. Set a small standing annual training budget per person for AI courses and exams, replacing case-by-case approvals that quietly discourage anyone from asking.
4. Ask every team member to automate one small task in their own job this quarter — the cheapest way to surface your natural builders.

What we asked

Your answer to each question, scored 0-4. The full maturity ladders are in the appendix.

K1. **How hands-on is your IT team with AI tools?** — 2 / 4

"Most have tried AI tools; no structured training"

K2. **How many of your IT staff hold current cloud or AI certifications?** — 1 / 4

"A few cloud certifications; no AI-related ones"

K3. **Does anyone on staff build with AI (APIs, automations, agents)?** — 1 / 4

"Individuals experiment on their own time"

K4. **How much time and budget do you invest in AI upskilling?** — 1 / 4

"Case-by-case approvals, no budget"

PART 06.5 · DIMENSION 5 OF 6 — 38 / 100

Use Cases & Value

Do you know where AI will pay off, and can you prove it?

A clear goal exists, but no prioritized list, no pilot, and no way to measure value.

You're ahead of most organizations on the question that matters — what is AI for? — because you have a specific, measurable target: cut service desk and documentation time. Your painful-processes list (password resets and access requests, SOPs and knowledge base articles, customer quotes) reads like a ranked backlog of AI candidates, and the first two are squarely in the sweet spot of what's proven today. What's missing is everything around the goal: no prioritized list with business owners, no pilot underway, and value judged by anecdote.

The measurement gap is the one that will hurt. You told us a win looks like a plan you can take to the CFO — but a plan without a baseline produces a pilot whose results are arguable. 'The team feels faster' does not survive a budget review; 'first response on our top ten ticket types is down 40% against a 90-day baseline' does. Pull the baseline before the pilot touches a single ticket. It's two weeks of work, it costs nothing, and it converts your pilot from an experiment into evidence.

Sequence the backlog deliberately: service desk first because it's measurable and internal, documentation second because it compounds (every SOP the assistant helps write improves the assistant's own source material), and customer quotes third because it touches pricing data — exactly the data you're cleaning up in the first 60 days — and deserves the guardrails being in place first.

STRENGTHS

- You have a concrete, measurable goal: reducing service desk and documentation time.

GAPS

- Ideas aren't prioritized or owned by the business.
- Value is judged by anecdote, so wins won't survive a budget review.

What to do about it

1. Baseline the service desk now: 90 days of ticket volume, first-response time, and time-to-resolve, tagged by your top 10 ticket types, plus a two-week sample of documentation hours.
2. Turn your painful-processes list into a ranked use-case register with a business owner and a success metric for each entry, reviewed by the AI council.
3. Scope the pilot narrowly: AI-drafted responses and knowledge articles for the top 10 ticket types, humans approving every send, judged on the two metrics leadership pre-agrees to.
4. Schedule quote preparation as use case three, gated on the sharing cleanup and DLP work being done, since it runs straight through your most sensitive data.

What we asked

Your answer to each question, scored 0-4. The full maturity ladders are in the appendix.

U1. **How well defined are your AI use cases?** — 2 / 4

"A list exists but isn't prioritized"

U2. **How far has AI gone beyond experimentation?** — 1 / 4

"Individuals use chat tools on their own"

U3. **What executive sponsorship and funding does AI have?** — 2 / 4

"Named executive sponsor, no dedicated budget"

U4. **How do you measure the value of AI?** — 1 / 4

"Anecdotes and user feedback"

PART 06.6 · DIMENSION 6 OF 6 — 25 / 100

Governance & Operating Model

Are rules, owners, and oversight in place before AI scales?

The weakest area. No policy, no decision owner, and no inventory of AI in use.

Governance is your lowest score and the root cause of most of the risk elsewhere in this report. Nobody owns AI decisions, so nobody can say yes safely — and when nobody can say yes, employees stop asking and adopt tools themselves, which is precisely the shadow-AI pattern in your Security score. The draft acceptable-use policy sitting unpublished is the purest example of cheap value waiting: it's mostly written, it costs nothing to finish, and it's the prerequisite for every other control you'll deploy this quarter.

Your regulatory surface makes the inventory gap concrete rather than theoretical. Under SOX, financial data flowing through an unvetted AI tool is an auditor's finding waiting to be written; under state privacy laws, you can't honor deletion or disclosure obligations for data you don't know left the building. An inventory — which tools, which data, which owner — is the two-week artifact that turns 'we think we know what's in use' into an answer you can give an auditor with a straight face.

Keep the machinery light. A company your size doesn't need an AI governance office; it needs one accountable council — IT, security, legal, finance, one business leader — meeting for thirty minutes every two weeks with a one-page intake form. The point is not process for its own sake; it's that a functioning council converts AI from a series of surprises into a queue of decisions.

STRENGTHS

- An AI acceptable use policy is already in draft.

GAPS

- IT decides on AI by default, without security, legal, or business input.
- SOX and state privacy obligations haven't been assessed against AI use.
- There's no inventory of AI tools or AI features inside your SaaS apps.

What to do about it

1. Publish the acceptable-use policy within 30 days: three data tiers (never / approved tools only / fine anywhere), named approved tools, a 10-minute training, and acknowledgment tracked in your existing policy tool.

2. Stand up the AI council with a standing 30-minute biweekly slot and a one-page intake form; its first two decisions are the approved assistant and the pilot scope.

3. Build the AI inventory from web-filter and SSO logs plus the AI features in your top 15 SaaS apps, and give every entry an owner and a data classification.

4. Have compliance map SOX and state-privacy obligations to AI use once, in writing, so the pilot starts with known constraints instead of discovered ones.

What we asked

Your answer to each question, scored 0-4. The full maturity ladders are in the appendix.

G1. Do you have an AI acceptable use policy? 1 / 4

"Drafting one"

G2. Who decides which AI tools and projects get approved? 1 / 4

"IT, by default"

G3. How have you addressed the AI regulations and standards that apply to you? 1 / 4

"Aware, but not assessed"

G4. Do you keep an inventory of AI in use, including AI features inside your SaaS apps? 1 / 4

"We could list the major tools if asked"

PART 07

Top risks

Red-flag answers first, then your lowest-scoring areas. Each with its concrete impact and the specific fix.

01 Overshared files become AI answers

Critical

IMPACT Any AI assistant connected to Microsoft 365 can surface files a user can technically open, such as pricing sheets, HR records, or board decks, even if they'd never find them on their own. Search friction is the only thing protecting that content today, and an assistant removes the friction completely. For a SOX-regulated manufacturer, this turns a permissions backlog into a live disclosure incident on day one of any Copilot-style rollout.

FIX Run SharePoint and OneDrive sharing reports, remove 'Everyone' and broad organization-wide links on your top 20 sites, and assign an owner to each site before enabling any Microsoft 365 AI assistant. Make 'sharing reviewed' a hard gate in the rollout plan rather than a parallel workstream, and put the remaining long tail of sites on a quarterly review cycle.

02 Customer and pricing data in public AI

Critical

IMPACT Your stated concern is already happening: employees are pasting work data into personal AI accounts with no controls. Data submitted to consumer AI tools can be retained under terms nobody has reviewed, which is a SOX and customer-contract problem, and you have no record of what has already left. Every week without an approved alternative deepens the habit and widens the exposure.

FIX Give everyone an approved enterprise AI assistant with data protection, then block or warn on unapproved AI sites through your web filter. The approved tool has to come first, or usage just moves to personal phones where you have no visibility. Pair the rollout with the acceptable-use policy so people know what's allowed, not just what's banned.

03 No inventory of AI in use

High

IMPACT AI features are switching on inside SaaS tools you already pay for, and employees are adopting tools on their own. Without an inventory you can't answer an auditor, a customer security questionnaire, or a breach investigation — and under state privacy laws you can't honor obligations for data flows you don't know exist. The gap compounds quietly with every renewal cycle.

FIX Pull a list of AI tools from web proxy and SSO logs, add the AI features in your top 15 SaaS apps, and keep it in a simple register with an owner and data classification for each entry. Route new entries through the AI council's one-page intake so the register stays current instead of becoming a snapshot.

04 No one owns data quality

High

IMPACT AI answering from your ERP and CRM will repeat today's data errors fluently and with confidence, which is worse than the current state: people trust a well-written wrong answer more than an obviously messy record. The first time the pilot confidently cites a stale price or a dead customer record, trust in the whole program takes the hit, not just the data.

FIX Name a data owner for the ERP, the CRM, and the service desk knowledge base, each responsible for a short list of quality rules for the fields the pilot will actually touch. Scope ruthlessly to what the pilot reads — this is about protecting the pilot's credibility, not launching a data governance program.

05 No AI credentials on the team

Medium

IMPACT Without a shared baseline, the team will struggle to evaluate vendors, secure new AI features arriving in every renewal, and build the service desk automation your goal requires. The practical cost lands in Days 61-90 of your own roadmap: the pilot needs a builder, and today you'd be hiring or renting one.

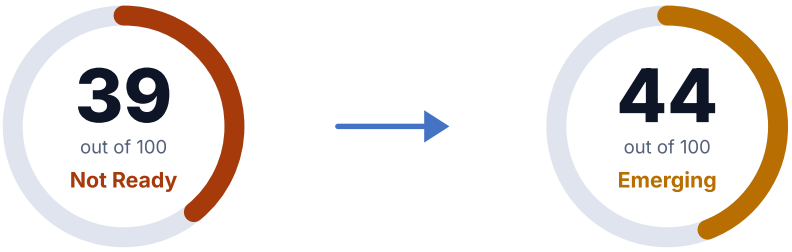
FIX Put every IT staff member through a fundamentals certification this quarter and send one engineer down an AI developer path on Azure now, so the skill exists in-house before the pilot build starts. Fund it from the training line — it's the cheapest risk reduction in this report.

PART 08

Quick wins

Low-scoring areas where the next maturity level is a low-effort fix. Weeks, not quarters.

WHAT THESE FIVE BUY YOU



Recomputed with the same scoring rules, assuming each quick win below moves up exactly one maturity level. No other improvement is assumed — this is the floor of what the next few weeks can do, not a projection.

EVERY LOW ANSWER, PLACED BY EFFORT TO FIX

Each chip is a question that scored 2 or below, by how many points it costs you and how hard it is to fix. The appendix has every question in full.

Points lost

	Low effort	Medium	High effort
-4 pts	DO FIRST		
-3 pts	K2 K4 U4 G1 G2 G4	D2 D3 S1 U2 G3	K3
-2 pts	S4 I4 K1 U1	S2 I2 I3 U3	LATER: D1 D4 I1

01 Start a team AI certification sprint

IT Director · Start this week; 6 weeks to first passes

1. Pick one fundamentals cert for the whole IT team (Azure AI Fundamentals fits your stack).
2. Give everyone 2 hours a week of protected study time for 6 weeks.
3. Track practice-test readiness weekly and book exams as people hit the bar.

Payoff: A shared AI vocabulary across the team and the in-house skills your pilot build depends on.

02 Set a small AI training budget

IT Director · 1 week

1. Set aside a per-person annual amount for AI training and exams.
2. Protect a recurring learning block on calendars.
3. Ask each person for one AI use they'll try in their own job this month.

Payoff: Removes the silent friction of case-by-case approvals and signals the team is being invested in.

03 Baseline your service desk before AI

Service Desk Manager · 2 weeks

1. Pull 90 days of ticket volume, first-response time, and time-to-resolve.
2. Tag the top 10 ticket types by volume.
3. Record time spent writing internal documentation.
4. Agree on the two metrics leadership will judge the pilot by.

Payoff: Turns the pilot from an arguable experiment into measurable evidence your CFO can fund.

04 Publish the AI acceptable use policy

IT Director with HR and Legal · 3 weeks

1. Finish the draft with three data tiers: never, only in approved tools, and fine anywhere.
2. Name the approved tools explicitly.
3. Have legal and HR review, then publish with a 10-minute training.
4. Require acknowledgment through your existing policy tool.

Payoff: Every employee knows what data can go where — the prerequisite for every other control this quarter.

05 Stand up a small AI council

Executive Sponsor · 2 weeks to first meeting

1. Invite IT, security, legal, finance, and one business leader.
2. Meet every two weeks for 30 minutes with a one-page intake form for new AI tools and ideas.
3. Make the first decisions: the approved assistant and the pilot use case.

Payoff: AI stops being a series of surprises and becomes a queue of owned decisions.

PART 09

The 90-day plan

Sequenced so risks close first, then foundations, then a pilot that proves value. Owners are roles, so you can assign names on day one.

DAYS 1-30
Close the risk gaps

DAYS 31-60
Build the foundations

DAYS 61-90
Prove value with a pilot

WHO IS LOADED WHEN

OWNER	DAYS 1-30	DAYS 31-60	DAYS 61-90
IT Director	●●	●	●
Security Lead	●	●	—
Microsoft 365 Admin	●	—	—
Service Desk Manager	●	●	●●
Executive Sponsor	—	●	—
Compliance Lead	—	●	—
Cloud Engineer	—	—	●

Days 1-30 — Close the risk gaps

Publish the AI acceptable use policy and run the 10-minute training. — IT Director
→ Every employee knows what data can go where.

Roll out an approved enterprise AI assistant and restrict unapproved AI sites. — Security Lead
→ Shadow AI drops and usage becomes visible.

Clean up sharing on the 20 highest-risk SharePoint sites. — Microsoft 365 Admin
→ Oversharing risk contained before any AI assistant is connected to company content.

Build the first AI inventory and stand up the AI council. — IT Director
→ One list and one decision-maker for AI.

Start the service desk baseline pull and the team certification sprint. — Service Desk Manager
→ Evidence and skills are building while the guardrails go in.

Days 31-60 — Build the foundations

Extend sensitivity labels and DLP to cover customer and pricing data. → Sensitive data is labeled and protected in email, endpoints, and AI tools.	Security Lead
Name data owners for ERP, CRM, and the knowledge base. → Clear accountability for the data AI will use.	Executive Sponsor
Map SOX and state privacy requirements to AI use. → Known obligations before the pilot touches production data.	Compliance Lead
Complete the team certification sprint. → A shared AI baseline across the IT team.	IT Director
Curate the service desk knowledge base and finalize pilot scope with the AI council. → The pilot will answer from trusted content against pre-agreed metrics.	Service Desk Manager

Days 61-90 — Prove value with a pilot

Launch a service desk pilot: AI-drafted ticket responses and knowledge articles for the top 10 ticket types. → A working pilot on real tickets with humans approving every response.	Service Desk Manager
Compare pilot metrics to the baseline every two weeks. → Measured time savings leadership can trust.	Service Desk Manager
Set Azure budgets and alerts for AI usage. → No surprise invoices as usage grows.	Cloud Engineer
Present results and a scale-or-stop decision to the AI council. → A funded next step based on evidence.	IT Director

PART 10

Months 4-12: the outlook

Where this program goes once the first 90 days have done their job.

The 90-day plan ends with a decision, not a finish line: the pilot's numbers go to the AI council, and what worked gets scaled while what didn't gets retired. The next nine months are about converting one proven use case into a repeatable pattern — same baseline-pilot-measure loop, applied down your painful-processes list — while the governance you stood up keeps pace with adoption.

MONTHS 4-6

Scale what the pilot proved

- Expand the service desk assistant from the pilot group to the full team, keeping human approval on customer-visible responses until quality data says otherwise.
- Point the documentation workflow at your SOP backlog: AI-drafted procedures reviewed by process owners, which compounds by improving the assistant's own source material.
- Make the Copilot decision company-wide using the finance trial readout and pilot data — expand, target specific roles, or consciously stop.
- Hold the first quarterly AI council review: refresh the inventory, review web-filter logs and DLP hits, and re-rank the use-case register.
- Report pilot ROI against the baseline to the CFO and lock funding for the next two use cases.

MONTHS 7-12

Systematize and extend

- Take on customer quote preparation — your third painful process — now that labels, DLP, and sharing cleanup protect the pricing data it runs through.
- Automate the top password-reset and access-request flows end to end, shifting the service desk assistant from drafting answers to resolving tickets.
- Graduate data ownership into a light governance rhythm: quarterly quality reviews for the systems AI reads, with metrics owners actually check.
- Decide the integration-layer question deliberately: whether a managed connector platform is justified by the use cases now in the register.
- Run the second skills wave: a second engineer on the Azure AI path and fundamentals for new hires, so the program isn't one person deep.
- Do a SOX-and-privacy readiness pass on everything now in production, so audit season documents a governed program instead of discovering an ungoverned one.

PART 11

Budget guidance

How to sequence the spend — and where not to spend at all.

Your stated range of $25K-$100K is comfortably enough for year one, and the plan is deliberately back-loaded: almost everything in the first 30 days is internal time rather than spend, the middle phase mostly activates capabilities in licensing you already own, and the meaningful money waits until the pilot has produced evidence. That sequencing is also your CFO story — by the time you ask for real spend, you're funding a measured expansion, not a bet.

ITEM	WHEN	SPEND	GUIDANCE
Acceptable-use policy, AI council, AI inventory	Days 1-30	No new spend	Internal time from IT, legal, and HR; the highest-leverage work in the whole plan.
SharePoint sharing cleanup	Days 1-30	No new spend	Admin time using reports already in your Microsoft 365 tenant; consider a one-time governance tool only if manual review stalls.
Approved enterprise AI assistant	Days 1-30	Included in existing licensing to Small	Start with entitlements in your current Microsoft agreement and the existing trial seats; expand only on usage data.
Sensitivity labels and DLP for customer and pricing data	Days 31-60	Included in existing licensing	Verify your Microsoft 365 plan tier covers the compliance features before buying add-ons.
Team certification sprint	Days 31-60	Small	Exam vouchers and study materials for the team, plus one deeper engineering path.
Service desk pilot (Azure usage)	Days 61-90	Small to Moderate	Usage-based spend; set the budgets and alerts before the first token, and treat the pilot bill as a data point.
Scale-up reserve	Months 4-12	The larger share of the budget	Hold in reserve for expanding what the pilot proves — full-team rollout, the documentation workflow, and use case three.

PART 12

Measuring success

The measures leadership should see, how to baseline them, and how often to look.

METRIC	BASELINE	TARGET	CADENCE
First-response time, top 10 ticket types	90-day pull from the ticketing system, tagged by type	Meaningfully faster on AI-assisted responses; agree the threshold with leadership up front	Biweekly during the pilot, monthly after
Time-to-resolve, top 10 ticket types	Same 90-day ticketing pull	Trending down without a quality regression	Monthly
Documentation hours per week	Two-week time sample across the IT team	Substantially less drafting time on AI-assisted SOPs and KB articles	Monthly
Unapproved AI usage	Web-filter log review before the approved assistant ships	Sustained decline toward zero after rollout	Weekly for the first 60 days, then monthly
Overshared sites remaining	Sharing reports on the top 20 sites	Zero broad links on the top 20, then extend to the long tail	Monthly until clean, then quarterly
Policy acknowledgment rate	Zero at publication	Everyone within 30 days of publishing	One-time, then quarterly for new hires
AI inventory coverage	First register from proxy and SSO logs plus top-15 SaaS review	Every discovered tool registered, approved, or blocked	Quarterly at the AI council
Certifications passed	Zero today	Whole-team fundamentals inside a quarter; one engineer on the deeper path	Monthly during the sprint

Baseline before the pilot starts. A measured pilot is evidence; an unmeasured one is an anecdote.

PART 13

Skills & certification plan

The gap between the plan in this report and the team that has to deliver it.

Your team needs a shared AI baseline now and one builder by the time the pilot starts. Because you're a Microsoft 365 and Azure shop, stay on the Microsoft path and add one vendor-neutral cert for the people who own risk.

ALL IT STAFF, INCLUDING THE SERVICE DESK

Microsoft Azure AI Fundamentals (AI-901)

A common vocabulary for AI on the platform you already run, in about six weeks of part-time study.

ONE SYSTEMS OR CLOUD ENGINEER

Microsoft Azure AI Apps & Agents Developer (AI-103)

Builds the skills to deliver the service desk pilot on Azure instead of waiting on a vendor.

IT MANAGER OR SECURITY LEAD

CompTIA AI Fundamentals

A vendor-neutral grounding for writing policy, evaluating AI vendors, and running the AI council.

Get your whole team exam-ready

HOW TO // AI CERT for Teams gives every person practice questions, timed mock exams, an AI tutor, and a readiness score for all of these certifications, with volume pricing.

See team pricing →

PART 14

A final word

You're closer than the score suggests. Most of the gap is decisions and configuration, not budget or new platforms. Get the guardrails in place this month, baseline before you build, and let the pilot's own numbers make the case for what comes next. Do that, and your team will be the one leading AI adoption, not cleaning up after it.

DD

David Davis

25+ years in enterprise IT across virtualization, cloud, and networking, now focused on AI. Author of 100+ courses on Pluralsight and LinkedIn Learning, co-founder of ActualTech Media, and CCIE #9369.

What to do with this report

Forward Parts 02-03 to your leadership team as-is. Assign a name to every owner role in Part 09 this week. Bring the budget table in Part 11 to your CFO alongside the baseline plan in Part 12 — that pairing is what turns this document into a funded program.

PART A

Appendix: your answers

Every question, the full 0-4 maturity ladder, and where you placed. Use the ladders as a target-setting tool: the level above your highlight is next quarter's goal.

A.1 — Data Readiness

38 / 100

D1. Where does most of your core business data live today?

0 Spreadsheets, file shares, and email attachments; no single source of truth

1 Several business apps (ERP, CRM, HR) that don't talk to each other

2 Core apps with some point-to-point integrations; reporting is still manual

3 A central warehouse or lakehouse covers some key domains

4 A governed central data platform covers most domains, with documented access

D2. Who is accountable for data quality?

0 Nobody; issues get fixed when someone complains

1 IT fixes problems ad hoc

2 Named owners exist for a few critical systems

3 Named owners for most data domains, with documented quality standards

4 A formal data governance program tracks quality metrics

D3. How well organized and permissioned is your unstructured content (SharePoint, Google Drive, file shares, wikis, ticket history)?

0 Sprawling; we don't know who can see what

1 We know there's oversharing but haven't addressed it

2 Permissions cleaned up on some high-risk sites

3 Most repositories have owners, retention rules, and reviewed permissions

4 Content is curated and labeled, and permissions are reviewed on a schedule

D4. Can AI tools get to your data programmatically?

- 0 No; data is exported by hand when needed
- 1 A few systems have APIs, but we don't use them
- **2 Some APIs or connectors in use, built one-off**
- 3 Most key systems are reachable through APIs or standard connectors
- 4 A managed integration layer with access controls covers our key systems

A.2 — Security & Privacy

50 / 100

S1. Which statement best describes employee use of AI tools today?

0 We don't know what AI tools people are using

1 People use public AI tools with work data; we have no controls yet

2 We block some public tools but offer no approved alternative

3 We provide approved AI tools and discourage or block unapproved ones

4 Approved enterprise AI tools for everyone; unapproved use is monitored and blocked

S2. How do you classify and protect sensitive data?

0 No classification scheme

1 A classification policy exists on paper only

2 Sensitivity labels applied to some data, mostly by hand

3 Labels plus DLP enforced on email, endpoints, and cloud storage

4 Labels and DLP also cover AI tools (prompts, uploads, and AI-generated content)

S3. How mature is your identity and access management?

0 MFA is not enforced everywhere

1 MFA everywhere, but access is broad and rarely reviewed

2 SSO for most apps; role-based access in key systems

3 SSO, role-based access, and periodic access reviews

4 Conditional access, least privilege, regular access reviews, and managed service/API identities

S4. Before adopting an AI tool or AI feature, do you review how the vendor handles your data?

0 No review

1 Ad hoc, when someone thinks to ask

2 Standard vendor security review, nothing AI-specific

3 Vendor review includes AI questions (training on our data, retention, data residency)

4 Formal AI risk assessment, including AI-specific threats such as prompt injection

A.3 — Infrastructure & Platforms

50 / 100

I1. Which best describes your infrastructure today?

0 Mostly on-premises, no cloud plans

1 Mostly on-premises, exploring cloud

2 Hybrid: key SaaS apps plus some cloud workloads

3 Majority cloud/SaaS with a defined cloud strategy

4 Cloud-first with a governed landing zone and infrastructure as code

I2. What access does your organization have to enterprise AI platforms?

0 None

1 Free or consumer AI tools only

2 Evaluating or piloting an enterprise AI assistant (e.g., Microsoft 365 Copilot, ChatGPT Enterprise, Gemini, Claude)

3 An enterprise AI assistant licensed for at least one department

4 Enterprise AI assistant broadly available, plus cloud AI services for building (e.g., Amazon Bedrock, Azure AI, Google Vertex AI)

I3. How automated are your IT operations?

0 Mostly manual; little scripting

1 Individual scripts, not shared or version-controlled

2 Shared scripts or an automation/RPA tool in limited use

3 Automation is standard: version control, runbooks, an integration platform

4 Infrastructure as code, CI/CD pipelines, and API-first integration are the norm

I4. Can you see and control what you spend on cloud and SaaS?

0 No visibility until the invoice arrives

1 Monthly invoice review only

2 Cost dashboards for major services

3 Budgets, alerts, and tagging for most services

4 Showback or chargeback to business units, with budget alerts and usage monitoring

A.4 — Skills & Talent

31 / 100

K1. How hands-on is your IT team with AI tools?

- 0 Little or no hands-on use
- 1 A few enthusiasts, self-taught
- **2 Most have tried AI tools; no structured training**
- 3 Most of the IT team has completed structured AI training
- 4 Ongoing, role-based AI training with hands-on labs

K2. How many of your IT staff hold current cloud or AI certifications?

- 0 None
- **1 A few cloud certifications; no AI-related ones**
- 2 Several cloud certifications; one or two AI-related
- 3 AI-related certifications held by people in key roles (cloud, security, data)
- 4 Certification paths are funded, tracked, and tied to role expectations

K3. Does anyone on staff build with AI (APIs, automations, agents)?

- 0 No
- **1 Individuals experiment on their own time**
- 2 A few people have built internal tools or automations
- 3 A named person or small team owns AI builds
- 4 A team builds with engineering discipline: testing, version control, monitoring

K4. How much time and budget do you invest in AI upskilling?

- 0 None
- **1 Case-by-case approvals, no budget**
- 2 General training budget, nothing AI-specific
- 3 Dedicated AI training budget and protected learning time
- 4 Dedicated budget, protected time, and an internal champions program that reaches business users

A.5 — Use Cases & Value

38 / 100

U1. How well defined are your AI use cases?

0 None identified

1 General ideas, nothing written down

2 A list exists but isn't prioritized

3 A prioritized backlog with business owners

4 Prioritized by value and feasibility, each with success metrics

U2. How far has AI gone beyond experimentation?

0 No sanctioned AI use

1 Individuals use chat tools on their own

2 One or two formal pilots underway

3 At least one use case in production with measured results

4 Multiple use cases in production and scaling

U3. What executive sponsorship and funding does AI have?

0 No executive interest

1 Interest, but no sponsor or budget

2 Named executive sponsor, no dedicated budget

3 Sponsor plus funded pilots

4 AI is in the strategic plan with multi-year funding

U4. How do you measure the value of AI?

0 We don't

1 Anecdotes and user feedback

2 Usage and adoption metrics (active users, seats used)

3 Business outcome metrics for some use cases (time saved, cost, quality)

4 Baselines captured before rollout; outcomes tracked and reported to leadership

A.6 — Governance & Operating Model

25 / 100

G1. Do you have an AI acceptable use policy?

0 No

1 Drafting one

2 Published, but no training or acknowledgment

3 Published, trained, and acknowledged by employees

4 Reviewed on a set schedule and backed by technical controls

G2. Who decides which AI tools and projects get approved?

0 Nobody; it happens ad hoc

1 IT, by default

2 An informal group, when issues come up

3 A cross-functional group (IT, security, legal, business)

4 A cross-functional AI council with a defined intake, review, and approval process

G3. How have you addressed the AI regulations and standards that apply to you?

0 Not considered

1 Aware, but not assessed

2 Legal or compliance consulted on specific projects

3 Applicable requirements mapped (e.g., EU AI Act, state AI laws, industry rules)

4 A formal framework adopted (e.g., NIST AI RMF or ISO/IEC 42001) with regular reviews

G4. Do you keep an inventory of AI in use, including AI features inside your SaaS apps?

0 No inventory

1 We could list the major tools if asked

2 Partial inventory, maintained by hand

3 Complete inventory with an owner and data types for each tool

4 Complete inventory with risk ratings, reviewed regularly, and AI covered in incident response

HOW TO AI

THIS IS WHAT YOU RECEIVE

Your report. Your stack. Your plan.

Every page of this document is rewritten from your own 24 answers and company context, reviewed personally by David Davis, and delivered as a private web report plus this PDF.

ready.howtoai.com

© 2026 HOW TO // AI

AI Readiness Report