SOLUTION BUILDERS®

A SOLUTION BUILDERS PLAYBOOK
2026 EDITION

A PLAYBOOK FOR LEADERS WHO WANT RESULTS, NOT DEMOS

The Practical AI Adoption Playbook.

A four factor score for picking use cases, a 30 day pilot design, the controls that keep company data out of public tools and a one page governance policy.

INSIDE THE PLAYBOOK

01 Use cases that pay off

02 A safe 30 day pilot

03 Data out of public tools

04 A one page AI policy

Published by Solution Builders
Bloomington, Minnesota · www.solutionbuilders.com

30+	Top 1%	98%
YEARS	OF MSPS	CSAT

WHAT'S INSIDE

Built for leaders who want AI to earn its keep.

Most companies are not short on AI ideas. They are short on a way to pick the right ones, prove them in a month and run them without leaking data. This playbook gives you all four pieces: a scoring model, a pilot design, a set of data controls and a policy people will actually read.

01 **Finding AI use cases that actually pay off.**
A four factor score that separates real productivity from demo theater, the use cases that pay back first and the ones to politely decline for now.
PAGE 03

02 **Piloting safely before you scale.**
How to design a 30 day pilot that proves or kills a use case: the group, the baseline, the metrics, the readout and the decision rule for day 30.
PAGE 06

03 **Keeping company data out of public tools.**
Shadow AI is already in your building. How data leaks through personal accounts, which tools do not train on your data and five controls that make the safe path the easy path.
PAGE 09

04 **Building an AI governance policy.**
The eight sections of a policy people will actually read, a data classification table you can lift as is, and the 90 day rhythm that keeps the policy alive.
PAGE 12

WHO IT IS FOR
Owners, COOs, CFOs and IT leaders at companies with 25 to 500 people who want AI results they can defend in front of a board.

HOW TO USE IT
Score your ideas with chapter 01, run one pilot with chapter 02, then put the guardrails from chapters 03 and 04 in place.

READING TIME
About 25 minutes cover to cover. The scoring sheet and the policy outline are built to copy into your own documents.

SOLUTION BUILDERS®

01

Finding AI use cases that actually pay off.

Every AI project that failed had a champion who believed in it. Belief is not a selection method. This chapter gives you a scoring model that ranks ideas on frequency, time, data and risk, so the first thing you build is the thing most likely to pay for itself.

65%
of organizations now use generative AI regularly in at least one business function

80%+
of AI projects fail, roughly twice the failure rate of ordinary IT projects

48%
of AI projects make it from prototype to production, a trip that averages eight months

30%
of generative AI projects will be abandoned after the proof of concept stage

Score every idea on four factors before you build anything.

Good AI use cases share a shape. The task happens often, it eats real time, the data it needs already exists in a usable form and a mistake is cheap to catch. Score each idea from 1 to 5 on the four factors below. The ranking usually makes the decision for you.

01 **Frequency**

How often does the task happen, and for how many people? A task 40 people do every day beats a task one person does once a month, however impressive the demo.

02 **Time per task**

How many minutes does it take today? AI pays back fastest on tasks that take 15 minutes to 2 hours: summarizing, drafting, searching, reformatting and first pass analysis.

03 **Data availability**

Does the information the AI needs already live in Microsoft 365, your ERP or your ticketing system? If step one is a six month data cleanup, score it low.

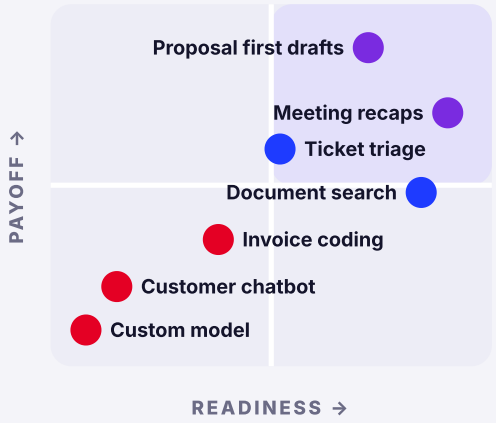
04 **Risk**

What happens when the output is wrong, and will a human see it before a customer does? Internal drafts score high. Anything that touches money, health or legal commitments scores low.

How to read the score

Add the four numbers. A 16 or above is a pilot candidate. A 12 to 15 goes on the backlog for when the data improves. Anything under 12 is a demo, not a use case, no matter who is excited about it.

Where common ideas land



Illustrative positioning from Solution Builders AI workshops. Readiness combines data availability and risk. Payoff combines frequency and time. Top right is where first pilots belong.

16+
points out of 20 makes an idea a pilot candidate

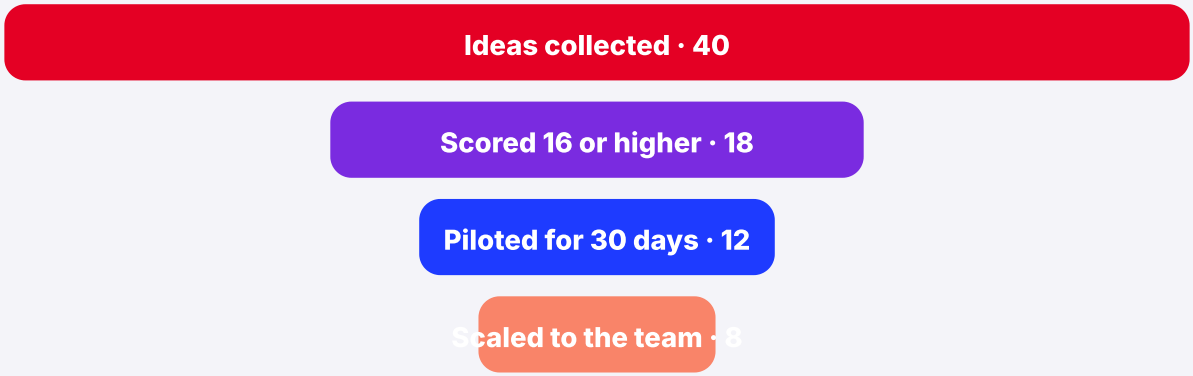
4
factors, each scored from 1 to 5, no weighting needed

> **The best first use case is one your team already does badly by hand, forty times a week.**
>
> SOLUTION BUILDERS AI PRACTICE

The best first use cases are boring, frequent and measurable.

Companies that see a return from AI in the first year almost never start with a customer facing chatbot. They start inside, on the work their people already complain about. Here is where the first payback shows up, and what to decline for now.

From a 40 idea workshop to production, over a year



Illustrative funnel from Solution Builders AI workshops. The drop from ideas to pilots is the scoring model doing its job, not a failure.

Use cases that pay back first

TEAM	USE CASE	WHAT TO MEASURE
Everyone	Meeting recaps with owners and actions	Minutes saved per meeting
Sales	First drafts of proposals and RFP answers	Hours per proposal, win rate held
Finance	Variance commentary and board pack drafts	Hours per close
Operations	Searching SOPs, contracts and past tickets	Time to answer, escalations avoided
Help desk	Ticket triage and suggested first replies	First response time, tickets per tech

Politely decline these for now

Fully automated customer replies, anything that makes a credit or hiring decision on its own, and training your own model. Each can work later. None is a first project for a company without an AI track record.

> **Pick the use case your CFO can measure without a spreadsheet full of assumptions.**
>
> SOLUTION BUILDERS AI PRACTICE

Sources: McKinsey, The State of AI in Early 2024; RAND Corporation, The Root Causes of Failure for AI Projects (2024); Gartner, 2024 survey on AI projects reaching production and 2024 forecast on generative AI project abandonment; Solution Builders AI workshop data.

02

Piloting safely before you scale.

A pilot is not a trial subscription with good intentions. It is an experiment with a start date, an end date, a baseline and a decision rule agreed in advance. Run it that way and 30 days is enough to know whether a use case deserves your money and your attention.

95%
of enterprise generative AI pilots show no measurable profit and loss impact

30
days is enough to prove or kill a use case, if the metrics are set on day one

12
people is a good pilot size: enough signal, small enough to support closely

11 min
saved per day is the point where people start to feel the tool is worth it

Design the pilot so it can fail cheaply and win clearly.

Most pilots drift because nobody wrote down what success looks like. Fix that before anyone gets a license. Four weeks, one use case, a dozen people, a baseline measured before the tool arrives and a decision meeting already on the calendar.

WEEK 1

Set up

Pick the use case and the 12 people. Measure the baseline on real work. Confirm data rules and enterprise licenses.

WEEK 2

Train and start

A 60 minute hands on session, a prompt library built for the task and a named channel for questions.

WEEK 3

Measure

Log time per task and run quality checks. A 15 minute midpoint check in fixes whatever is slowing people down.

WEEK 4

Decide

Compare to baseline, read the survey, price the rollout and make the scale, extend or stop call.

Before day one

- [x] One use case, written as a sentence: "Draft first versions of customer proposals from our template and past wins."
- [x] A baseline: how long the task takes today, timed on at least 20 real examples, not estimated.
- [x] A pilot group of 8 to 15 people who do the task weekly, including at least one skeptic.
- [x] Enterprise licenses only. No personal accounts, no free tiers, no data outside the approved tools.
- [x] A sponsor who will attend the day 30 readout, and a decision rule everyone has seen in writing.

Why the skeptic matters

Enthusiasts make every tool look good. The skeptic finds the prompts that fail, the outputs that need rework and the step nobody documented. If the tool wins them over, the rollout will go fine.

Pilot success metrics

METRIC	HOW YOU GET IT	BAR TO CLEAR
Time per task	Timed samples, before and during	**25% faster or better**
Quality	Manager spot checks on 10 outputs a week	**Equal to baseline or better**
Adoption	Usage logs, three sessions a week per person	**80% of the group active**
Rework	Outputs rewritten from scratch	**Fewer than 1 in 5**
Data incidents	DLP alerts and self reports	**Zero**

> **A pilot without a baseline is a free trial with a slide deck at the end.**
>
> SOLUTION BUILDERS AI PRACTICE

Measure hours saved the way a CFO would.

Self reported savings run high, so count only time you can see and redeploy. Here is a readout from a 12 person pilot, the math that turns minutes into a business case and the rule for day 30.

Minutes saved per person per day, by task



Illustrative readout from a 12 person pilot, measured with timed samples, not a survey.

The math

Minutes per day, times working days, times loaded hourly cost, minus licenses and the time to run the program. Then cut the result by a quarter for optimism. If it still clears your hurdle rate, scale it.

The day 30 decision

RESULT	WHAT YOU SAW	WHAT YOU DO
Scale	Beat the time bar, quality held, 80% of the group active	Roll out to the full team within 30 days, with training
Extend	Time saved, but adoption or quality wobbled	Fix the prompt library or the training, run 30 more days
Stop	Under 10% faster, or quality slipped	Write a one page lessons note and pick the next use case

46 min
saved per person per day across the five tasks

3.8 hrs
per person per week back for billable or strategic work

A DEMO

What vendors show you

- A perfect prompt on clean data.
- One enthusiastic user, no skeptics.
- No baseline, so everything looks like a win.

A PILOT

What you should run

- Real work, messy data, a skeptic in the group.
- A baseline measured before the tool arrives.
- A decision rule written down on day one.

Sources: MIT NANDA, The GenAI Divide: State of AI in Business 2025; Microsoft WorkLab, AI Data Drop: The 11 by 11 Tipping Point (2024); pilot readout figures are illustrative and drawn from Solution Builders client engagements.

SOLUTION BUILDERS®

03

Keeping company data out of public tools.

Blocking AI does not work. People route around blocks with their phones and personal accounts, and the company loses both the productivity and the visibility. The answer is to give people an approved tool that is better than the free one, then put quiet controls behind it.

78%
of people using AI at work bring their own tools, usually without telling IT

72%
of generative AI use on company devices runs through personal accounts

27%
of the data employees put into AI tools is sensitive, up from 11% a year earlier

$670K
added to the average breach cost when shadow AI is involved

Your people are already using AI. The question is where.

The risk is not the technology, it is the account. A free chatbot on a personal login may keep prompts, use them for training and sits outside every control you own. The same model under an enterprise agreement, or inside your Microsoft 365 tenant, is a different product.

Four tiers of AI tools

TIER	EXAMPLES	YOUR DATA	ALLOWED DATA
Free consumer	Personal ChatGPT, Gemini or Claude accounts	**May be kept and used for training**	**Public information only**
Paid consumer	Individual paid plans on a personal email	**Opt outs vary, no admin visibility**	**Public and low risk internal**
Enterprise	ChatGPT Enterprise, Claude for Work, Gemini for Workspace	**Not used for training, by contract**	**Internal and confidential, with DLP**
Inside your tenant	Microsoft 365 Copilot and Copilot Chat	**Stays in the tenant, honors permissions**	**Whatever the user can already open**

Three questions for any AI vendor

- **Do you train on our prompts or files?** The answer must be no, in the contract, not the FAQ.
- **Where is our data stored, and for how long?** A named region and a retention period you control.
- **Can our admins see usage and switch it off?** Single sign on, audit logs and central controls.

0

visibility your IT team has into prompts sent from personal accounts

1 in 5

organizations reported a breach involving shadow AI

The lesson from Samsung

In 2023 engineers at Samsung pasted confidential source code and meeting notes into a public chatbot while fixing a bug. The company found out afterward and restricted the tools company wide. Nobody involved was trying to leak anything.

> **People do not leak data because they are careless. They leak it because the safe option was slower.**
>
> SOLUTION BUILDERS AI PRACTICE

Five controls that make the safe path the easy path.

None of these require a new security team. Most are settings in Microsoft 365 and Entra ID that you already pay for, plus a short list of decisions. Done in this order, they take about two weeks.

1. **Publish the approved list**
   One page: which tools are approved, for which data class, and how to request a new one. Review it every quarter.

2. **Buy enterprise seats where needed**
   Copilot or an enterprise plan for the roles with real use cases. A paid seat is cheaper than one incident.

3. **Turn on data loss prevention**
   Purview DLP, included with Business Premium and E3, warns or blocks when confidential data heads for an unapproved site. Endpoint DLP needs E5 or an add on.

4. **Govern the long tail**
   Defender for Cloud Apps shows every AI app in use. Sanction the good ones, block the rest and tell people why.

5. **Give everyone a safe default**
   Copilot Chat with enterprise data protection comes with Microsoft 365. Make it the first thing people find when they search for "AI".

DLP rules worth turning on first

- [x] Block uploads of files labeled Confidential or Restricted to unapproved AI domains.
- [x] Warn on pastes that match card numbers, Social Security numbers or customer record patterns.
- [x] Block sign ins to consumer AI services with personal accounts on company devices.
- [x] Alert the security team when one user sends unusual volumes to any AI service.
- [x] Review the alert report monthly with the policy owner and tune the rules.

> **Block the tool and you lose the visibility. Approve a better one and you get it back.**
>
> SOLUTION BUILDERS AI PRACTICE

PUBLIC CHATBOT

What you get for free

- Prompts may be retained and used for training.
- No admin visibility, no logs, no off switch.
- Nothing it says is grounded in your files.

MICROSOFT 365 COPILOT

What the tenant version adds

- Prompts and files stay inside your tenant and never train the models.
- Answers grounded in your mail, files and meetings, limited to what each user may see.
- Audit logs, retention, DLP and sensitivity labels apply automatically.

Sources: Microsoft Work Trend Index 2024; Verizon 2025 Data Breach Investigations Report; Cyberhaven, AI Adoption and Risk Report 2024; IBM Cost of a Data Breach Report 2025; Microsoft documentation on enterprise data protection in Copilot.

04

Building an AI governance policy.

The best AI policy in your industry is one page long and everyone can quote it. The worst is 30 pages, written by counsel, read by nobody and discovered during an incident. This chapter gives you the short version, the table that does most of the work and the review rhythm.

63%
of breached organizations had no AI governance policy or were still writing one

60%
of leaders worry their organization lacks a plan and vision for AI

2
public frameworks to borrow from: the NIST AI Risk Management Framework and ISO/IEC 42001

1 page
is the right length for the policy. Put the detail in an appendix nobody has to memorize

A policy people read fits on one page.

A working AI policy answers eight questions in plain language and points to a table for the details. Here is the outline Solution Builders uses with clients, and the data classification table that does most of the work.

The eight sections

SECTION	WHAT IT SAYS, IN ONE LINE
Purpose and scope	Who the policy covers, including contractors, and which tools it applies to
Approved tools	The current list, by data class, and how to request an addition
Data classes	Public, internal, confidential and restricted, and what may go where
Acceptable use	What AI may draft, summarize and analyze, and what it may never decide alone
Human in the loop	A named person reviews anything that leaves the company or affects a person
Transparency	When to tell customers, candidates or regulators that AI was used
Security and privacy	Enterprise accounts only, DLP on, no personal data in prompts without a lawful basis
Review and incidents	Who owns the policy, how to report a slip without blame, and the quarterly review date

Data classes and where they may go

DATA CLASS	PUBLIC CHATBOT	ENTERPRISE AI	COPILOT IN TENANT
Public	Yes	Yes	Yes
Internal	No	Yes	Yes
Confidential	No	With DLP and labels	Yes
Restricted	No	No	Labeled and limited

Restricted means regulated or contractually protected data: health records, card data, CUI, source code under NDA.

Human in the loop, defined

A human reviews, and is accountable for, any AI output that goes to a customer, makes a decision about a person or commits the company to anything. Write the names of those humans into the policy.

> **If your people cannot quote the policy, you do not have one. You have a document.**
>
> SOLUTION BUILDERS AI PRACTICE

Governance is a rhythm, not a document.

A policy nobody revisits is out of date in a quarter, because the tools change monthly. Put four dates on the calendar, give the policy an owner and score yourself against the ten items below twice a year.

DAY 1

Publish and train

One page policy, approved tool list and a 30 minute session for every employee.

DAY 30

First review

Read the DLP and usage reports. Approve or decline tool requests. Fix the top confusion.

DAY 60

Pilot readouts

Scale, extend or stop decisions for the first use cases. Update the use case register.

DAY 90

Leadership review

Hours saved, incidents, spend and the next wave. Refresh the policy and repeat quarterly.

Governance checklist, part one

- A one page policy exists, is signed by leadership and was acknowledged by every employee.
- An approved tool list is published and reviewed at least quarterly.
- Data classes are defined and the "what goes where" table is in the policy.
- Every approved tool runs on enterprise accounts with single sign on and audit logs.
- Human review is required, and assigned by name, for customer facing output and people decisions.

Governance checklist, part two

- A named owner holds the policy, the tool list and the incident log.
- A use case register tracks every pilot with its score, owner, metrics and decision.
- Vendor contracts state that your data is not used for training and name a retention period.
- Employees know how to report an AI slip without fear, and at least one has done so.
- The policy was reviewed in the last six months against new tools and new regulations.

_____ / 10

Score yourself

Count the items fully in place and write the number here. Repeat every six months.

ITEMS DONE	WHERE YOU STAND
0 to 3	Shadow AI is running your AI program. Start with the policy and the tool list this month.
4 to 7	Typical. The gaps are usually vendor contracts, named owners and the review rhythm.
8 to 10	Governed. Spend the energy on finding the next use case.

Sources: IBM Cost of a Data Breach Report 2025; Microsoft Work Trend Index 2024; NIST AI Risk Management Framework 1.0 (2023); ISO/IEC 42001:2023; Solution Builders AI governance methodology.

AI that earns its keep, with the guardrails built in.

Solution Builders has run technology for Minnesota businesses since 1995. Our AI consulting practice sits inside the RISE framework, delivered by 50+ engineers, backed by a 24/7 security operations center and SOC 2 Type II audited.

R

Resilient

Cloud foundations and tested backups, so AI runs on infrastructure that stays up.

I

Integrated

AI connected to Microsoft 365, your line of business systems and clean data.

S

Secure

Data loss prevention, labels, enterprise accounts and a governed tenant.

E

Evolving

Use case scoring, pilots, Copilot readiness and a policy that keeps pace.

01

AI Opportunity Snapshot

A 30 minute conversation. You leave with your top three scored use cases and the one data risk to fix first.

FREE · 30 MINUTES

02

AI Readiness Assessment

Two weeks covering use cases, data, tooling and governance, with a scored report and a prioritized 90 day plan.

FIXED FEE · 2 WEEKS

03

AI Consulting Program

Everything in this playbook, run with you: pilots, Copilot rollout, controls, training and quarterly reporting.

MONTHLY · ONGOING

See your top three AI use cases in 30 minutes.

Book a free AI Opportunity Snapshot with a Solution Builders solution architect. No demo, just a scored list and a clear next step.

Book my AI snapshot

THE SHORT VERSION

Three things to remember.

If this playbook ends up in a drawer, keep these three ideas. They separate companies that get a return from AI from companies that get a subscription.

01 **Score before you build.**

Frequency, time per task, data availability and risk. A 16 or higher is a pilot. Everything else is a demo.

02 **Thirty days, one use case, a baseline.**

Measure before the tool arrives, decide on day 30 and count only the hours you can redeploy.

03 **Make the safe path the easy path.**

Enterprise accounts, DLP and a one page policy. People leak data when the approved option is slower.

SOLUTION BUILDERS®

START THE CONVERSATION

www.solutionbuilders.com

(952) 854-8188 · 1550 American Blvd East, Suite 900, Bloomington, MN 55425
Your technology partner since 1995

Book a free assessment

© 2026 Solution Builders, Inc. All rights reserved. Product names and marks are the property of their respective owners. Figures marked illustrative are planning estimates, not guarantees.