Praxent

Select Services Partner
Claude Partner Network

Microsoft Solutions Partner
Cloud & AI Platforms

2027 PLAYBOOK · BANKING

The 2027 AI Planning Playbook for *Banking* Market Leaders

10 Pillars for building an AI plan that delivers measurable results.



Your 2027 AI plan decides whether your bank *leads* the market

Like most banks, you've likely already bought AI. Now you're looking at 2027 to be the year it starts to pay off in your P&L, in front of your board, and in the exam room.

AI is already changing how banks compete, and even S&P believes we'll know which banks will lead within the next two to four years.[1]

Even after paying for AI, banks on average see a 10% to 25% increase in their efficiency. Those gains show up through lower costs, new revenue, and reduced risk.[1]

But getting there is still the largest hurdle. Moving AI from pilots into production remains the step where most financial services firms abandon nearly half of their AI projects.[2]

10% to 25%

average efficiency gain for banks, net of their AI investment

S&P Global, 2025

The banks that make this move successfully are the ones treating AI as an overall transformation of how their bank works.

Your current differentiators will serve as the basis for your next steps

Your people already know how your bank lends, serves, and stays compliant. Meanwhile, your core holds decades of the rules that have made it work well up until now.

A plan that incorporates both while adapting them for AI is how a careful bank becomes an AI-native one.

Only 8% of banks have AI fully integrated across their organization. Data points to three causes that every bank leader will recognize, from legacy technology to heavy regulation and the need for robust model-risk controls.[3]

Which is why a market leading AI plan for banking builds on all three.



8%

of banks have AI fully integrated across their organization

S&P Global Ratings, 2026

Build your *market-leading* AI plan around your core systems, regulators, and risk culture

01

Your core's decades of business rules

The plan recovers them first, so AI works with the logic that already runs your lending, deposits, and payments.

02

Your regulators that expect you to show your work

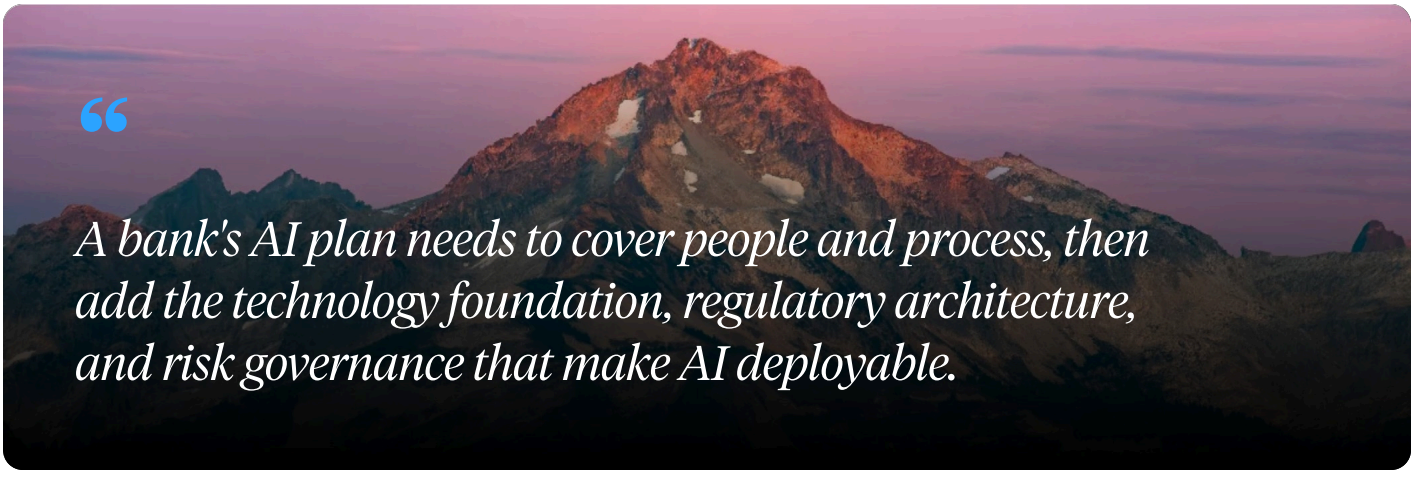
The plan puts validation, monitoring, and audit trails around every use of AI, so the answer is ready before an examiner asks.

03

Your internal risk culture as an asset

The plan seats risk, compliance, and audit on day one, and their caution turns into speed because nothing has to be rebuilt to pass review.

Most AI advice reads these three factors as the reasons banks move slowly, yet we see them as the foundation for **your successful AI transformation**.



Here are the 10 Pillars to build the AI plan that delivers measurable results.

NEXT PAGE →

The 10 *pillars* of a bank's AI plan



PHASE 1

Build the foundation

Settle what your board and your examiners will ask about before any AI goes live.

PILLAR 1

Strategic foundation

Decide what your AI is for, then tie every AI dollar to a result your board already tracks.

p. 05

PILLAR 2

Leaders and governance

Name five AI owners who keep decisions moving and traceable.

p. 08

PILLAR 3

Regulatory and risk

Map every rule to a control before AI reaches customers.

p. 10

PILLAR 4

Data readiness

Recover your core's rules and determine which data is AI ready.

p. 13



PHASE 2

Map your workflows

Pick the workflows that will pay for the rest of your AI plan.

PILLAR 5

The AI-native bank matrix

Start where the impact is high and effort is low, and let early wins fund the bigger bets.

p. 16



PHASE 3

Execute your plan

Put the plan to work, prove the return against your baseline, and show your board results.

PILLAR 6

The 30 / 90 / 365 roadmap

Open each phase only after the last one proves value that you can show your board.

p. 17

PILLAR 7

Upskilling

Train managers first as they set each team's pace, while you move every role up a proficiency level.

p. 19

PILLAR 8

ROI measurement and budget

Prove the return to the board in your own numbers, then fund AI as a standing budget line.

p. 21

PILLAR 9

Agent strategy

Hand agents the routine work so your people spend their day on judgment calls.

p. 24

PILLAR 10

Change management

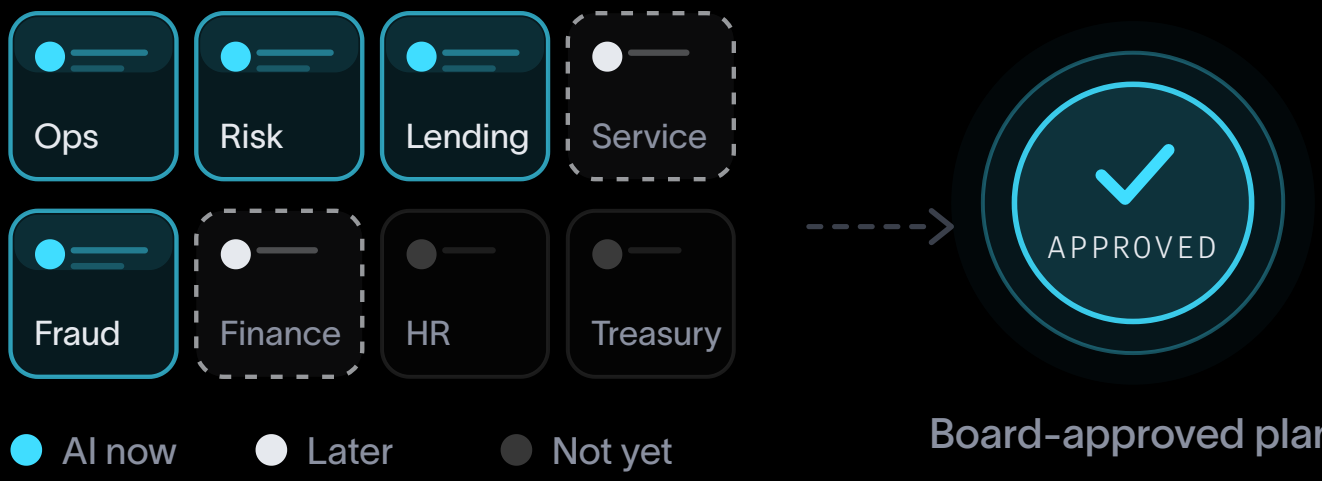
Close the four gaps between AI access and daily use to drive the greatest gains.

p. 27

PILLAR 1

Your strategic foundation

A board-approved plan that ties AI spending to results you already track



Over half of banks don't have an enterprise-wide AI strategy.[3] Of the ones who do, many only cover a handful business lines, with the rest still in draft.

Before you buy anything else, document what AI is for at your bank and specifically name the use cases that it will support.

Starting that one document puts your bank among the market leaders who have an enterprise-wide AI strategy.



56%

of banks lack an enterprise-wide AI strategy

S&P Global Ratings, 2026

Seven decisions your AI manifesto puts into writing

THE MANIFESTO

That document becomes your AI manifesto. It also gives your bank a board-approved roadmap that your examiners can read.

Your leadership should make the following seven decisions first and then revisit them each budget cycle, so no business line ever makes them alone.

YOUR LEADERSHIP TEAM DECIDES	WHAT THAT LOOKS LIKE IN BANKING
01 **Internal operations** — Where AI works inside the bank first	Back-office, compliance, and risk teams use AI to review loan documents and answer policy questions, which lowers cost-to-serve
02 **Customer-facing position** — Which products and services AI will change for customers	AI-native products, such as account opening a customer finishes on their phone in one sitting
03 **Internal boundaries** — What AI can and can't touch inside the bank	AI drafts the credit memo, and a loan officer makes the lending decision
04 **External boundaries** — What you owe customers and regulators when AI is involved	Disclosure when AI is used, specific reasons for any decline a model informed, and the decisions you keep with a person
05 **Cultural norms** — How employees work with AI and how managers set expectations	Approved AI is the norm for everyday work, and managers say where they expect it and how output gets checked
06 **Risk treatment** — How the bank evaluates, monitors, and governs AI risk	Every use case gets a risk tier and a named owner before launch, through the risk program you already run
07 **Budget allocation** — What the AI budget covers and who pays for what	Licenses, training, and process redesign in one budget, with platform costs shared by the business lines that use them

Tie AI to the goals your board already tracks

Tie your manifesto to goals the board already tracks. Then set how often AI appears on the board agenda, and report the same measures each time.

GOALS THE BOARD ALREADY TRACKS

- Deposit growth
- Loan volume
- Efficiency ratio

AI cuts your cost-to-serve and opens up revenue paths that were previously out of reach

LOWER COST-TO-SERVE

Fewer manual steps

AI takes the manual steps out of back-office, compliance, and risk work, so those teams spend their judgment on the calls that need it.

NEW REVENUE

Expanding services to customers

It also opens revenue you couldn't serve profitably before, such as small-business loans too small to underwrite by hand and personal guidance for customers who've never had a banker.

Your manifesto needs to name the results you expect on each side, so each of your bank's AI investments traces back to one of them.

HOW PRAXENT HELPS

Praxent's banking team maps your top AI opportunities to cost-to-serve and revenue, then helps your leadership team turn them into a manifesto your board can approve, with your first use cases tied to numbers the board already tracks.

Write your AI manifesto →

PILLAR 2

Leadership and governance

Name five execs as AI owners so decisions don’t stall between teams.



When every AI decision has a specific owner, your bank can move faster with less risk. Naming these owners before your first use case goes live keeps approvals moving forward.

Without board oversight, AI ownership scatters. Only 15% of institutions give it to the CEO and strategy teams.



39%

Institutions that include board-level oversight in their AI framework

S&P Global Ratings, 2026

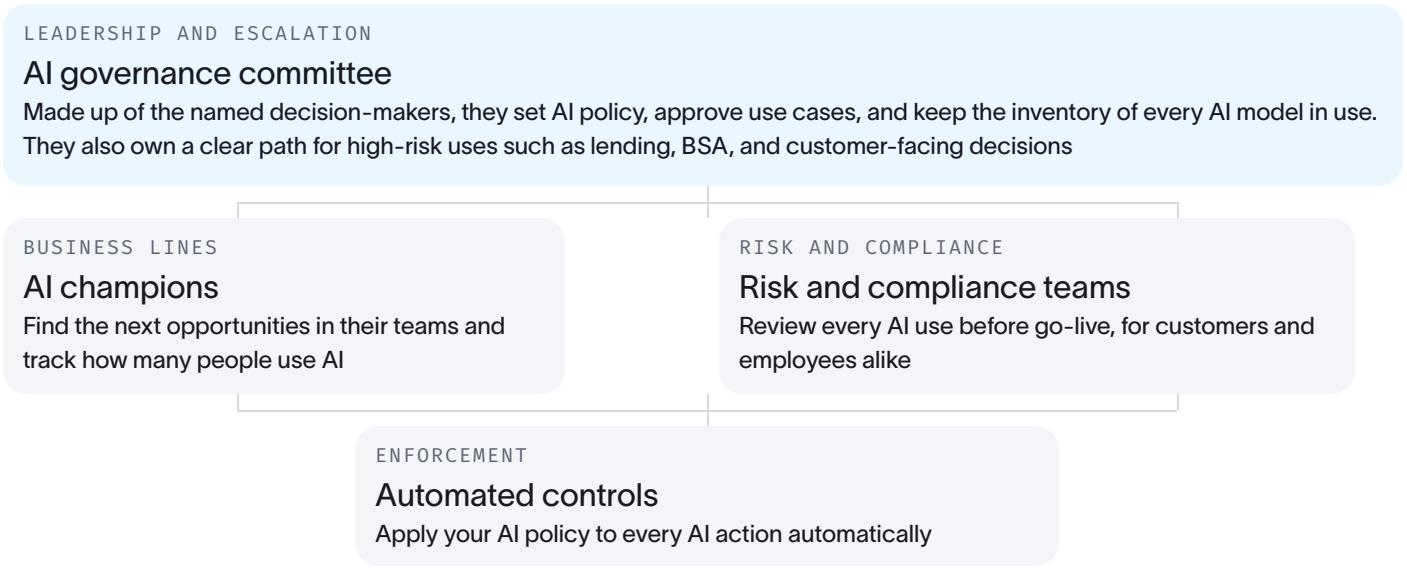
Five executives should own AI at your bank, led by a CEO who mandates adoption and reports to the board

SEAT	OWNER	WHAT THEY OWN
EXECUTIVE SPONSOR	CEO or president	Mandates AI adoption and reports to the board
PLATFORM AND SECURITY	CTO or CIO	The approved AI platform, its fit with your systems, and cybersecurity
DELIVERY	Head of AI	Daily execution, adoption numbers, and coordination across business, risk, and technology
RISK	Chief risk officer	Model risk, regulatory alignment, and the examiner relationship
COMPLIANCE	Compliance officer	Policy enforcement, fair lending, and BSA/AML validation

Together this leadership team forms one committee that sets AI policy and approves use cases. Each business line then names an AI champion who brings forward the opportunities their team sees.

Risk and compliance review every AI use before it goes live, and automated controls apply your policy to every AI action. That's how one policy holds across the bank without relying on anyone's memory.

THE GOVERNANCE MODEL



HOW PRAXENT HELPS

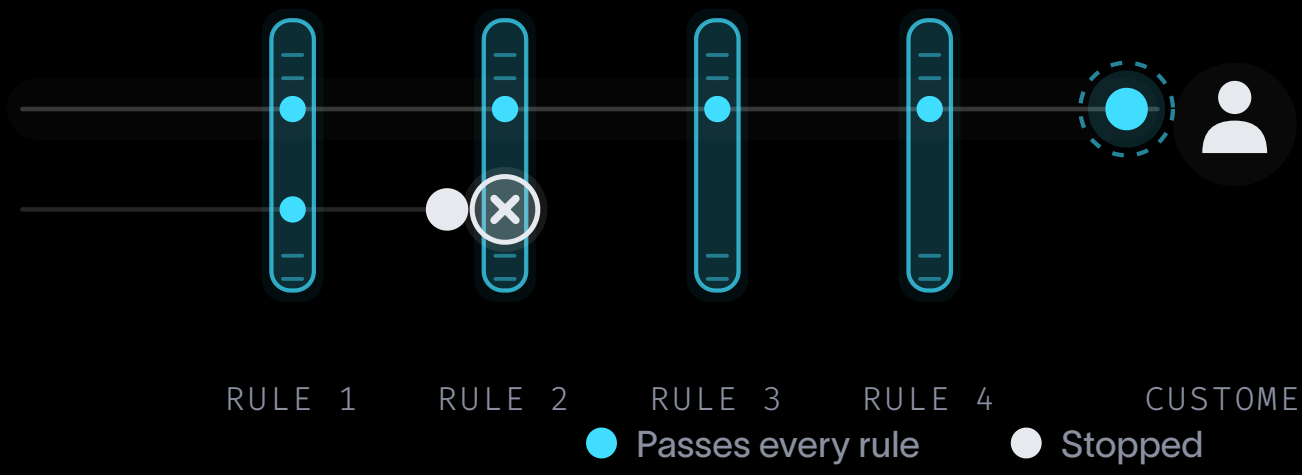
PraxentGuard℠, our AI controls framework, enforces your AI policy automatically on the AI work we deliver for your bank, so the policy your committee sets is the one every system follows.

Name your AI owners →

PILLAR 3

Regulatory and risk architecture

Set AI rules your examiners can check long before a customer ever sees them.



The standard for generative and agentic AI is yours to set.

In April 2026, SR 11-7 and OCC Bulletin 2011-12 were replaced by SR 26-2 and OCC Bulletin 2026-13. This new guidance purposefully left generative AI and agents out of its scope.

In other words, you get to set the standard.[4, 5] Put it in place now to avoid use cases stalling in review.



77%

of financial institutions named data privacy, security risk, and regulatory compliance as key challenges to them adopting AI

S&P Global Ratings, 2026

Map your AI to the four rule sets your examiners already use

For each rule, write down how your bank meets it. Be sure to recheck your map every time you update your plan, since the agencies have said more guidance is coming.

01

Model risk

SR 26-2 and OCC Bulletin 2026-13, April 2026 [4][5]

- Validation, monitoring, and a model inventory scaled to each model's materiality, plus your own standard for generative and agentic AI

02

Third-party risk

The 2023 interagency guidance, OCC Bulletin 2023-17 [8]

- Diligence scaled to each AI provider's risk, a view of how much you depend on any one of them, and an exit plan

03

Fair lending

ECOA and Regulation B, amended July 2026, plus the Fair Housing Act and state law [9][10][11]

- Specific reasons for any adverse action a model informed
- Outcome testing where the Fair Housing Act and states such as New York and Illinois still apply disparate impact

04

Data privacy and security

GLBA, the interagency information security standards, and the FFIEC IT Examination Handbook

- Customer data classified and inventoried before AI touches it
- AI tools held to the same cybersecurity standards as every other system

Classify every type of sensitive bank data with a clear AI access level

Start with the seven types of sensitive data that every bank holds. Make these decisions before deployment, so no team has to make that call in the middle of a project.

DATA TYPE	Never No AI access under any condition	With approval Governed use, with explicit sign-off for each use case	Allowed Cleared for standard workflows
01 Cardholder data	●	·	·
02 Consumer nonpublic personal information	·	●	·
03 KYC records and other personal information	·	●	·
04 Material nonpublic information	●	·	·
05 BSA/AML records	·	●	·
06 Credentials	●	·	·
07 Privileged legal material	·	●	·

This map is illustrative, and your risk team sets the final one.

Keep your AI records in one file so they're ready for your next exam

Examiners are already asking banks how limit, monitor, audit, and shut down AI agents.[12] Organize your AI policy, model inventory, test results, validation results, data classifications, vendor diligence, and approvals together in your AI records file. You'll stay ready for any questions examiners have.

HOW PRAXENT HELPS

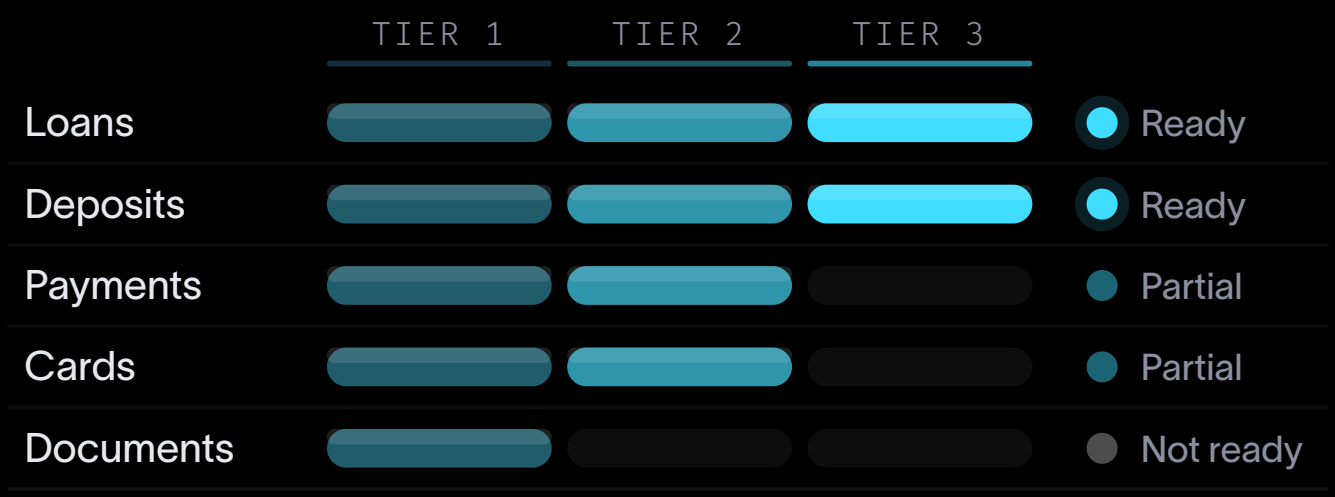
PraxentGuard℠ then enforces the boundaries you set on the AI systems Praxent delivers, so data marked "Never" stays out of reach.

Set your AI rules →

PILLAR 4

Technology and data readiness

Prepare your core, its rules, and your data for AI.



AI is only as useful as the data underneath it and financial institutions regularly cite data readiness as a key hurdle. An agent reviewing a BSA alert only helps your investigator when it can see the full transaction history, KYC file, and case notes.

Before you fund anything further, make sure your data is ready for AI and can serve as one source of truth the whole bank decides from.



35%

of the current AI workload at financial services firms goes to getting data ready for use

Three tiers of data readiness move your bank to AI that scales across every business line

Know which tier of readiness each part of your bank has reached before you fund the use cases that depend on it.

TIER 1

Foundational

WHAT'S IN PLACE

Centralized, governed access to data your people trust

WHAT IT LETS BANKS DO

Reports and answers your board and examiners can rely on, from one version of each number

TIER 2

Structural

WHAT'S IN PLACE

Systems that share data as it changes, with the same definitions and clear labels everywhere

WHAT IT LETS BANKS DO

A full, current view of each customer across lending, deposits, and servicing, for staff and agents alike

TIER 3

Scalable

WHAT'S IN PLACE

A record of every change to your data, built to grow from one pilot to the whole bank

WHAT IT LETS BANKS DO

Take what works in one business line to every line without starting over, and show where each number comes from

The result becomes one set of numbers the whole bank decides from.

Scalable data readiness means you can have a customer view that carries the full relationship with agents that follow the same rules your people do. Getting there starts with recovering the decades of business rules your core runs on and mapping which systems feed which processes.

This identifies which processes AI can reach today and which need modernizing first. Modernizing a system AI can't reach yet opens that process to AI across the whole business line, so put that work in the plan and the budget.

The six steps to put approved AI in people's hands across your bank, one team at a time

Before bank-wide rollout, list every core system AI will work with, what it can see and change in each, and which to connect first, and then rank them by value.

1

Readiness check

Confirm licenses, data access, and security reviews

→ The rollout won't stall halfway on a missing license or review

2

Bank-wide setup

Set up one approved AI platform for the whole bank

→ Every team works under the same terms and controls

3

Employee access

Connect the platform to the logins your people already use

→ Access follows each person's role and ends when they leave

4

Security settings

Apply your data boundaries and security settings

→ Your Pillar 3 decisions hold before anyone starts

5

Monitoring

Track usage, cost, and results from launch

→ You report adoption and return to the board with your own numbers

6

Team-by-team rollout

Open access in waves, starting with your first teams

→ Each wave learns from the last, and support stays manageable

HOW PRAXENT HELPS

PraxentDecode℠ recovers the business rules and dependencies inside your core, so the logic your bank runs on is documented before anything changes and you can modernize around it with nothing lost.

Your teams start using the approved AI platform before the engagement ends.

Check your data tiers →

PILLAR 5

The AI-native bank matrix

Find the workflows worth funding first

P&L IMPACT →

QUICK WINS	STRATEGIC BETS
High impact, low effort	**High impact, high effort**
Start here, and let these results fund the bigger bets.	Fund these once the quick wins pay off and the data is ready.
FOUNDATIONAL	**SEQUENCE CAREFULLY**
Lower impact, low effort	**Lower impact, high effort**
Workflows that build skills and clean up data along the way.	Schedule these behind the work that unlocks them.

TRANSFORMATION EFFORT →

Our AI-native matrix scores banking workflows across eight domains by P&L impact and effort. From consumer and commercial lending to payments, fraud, servicing, and compliance, you can find and fund the ones most relevant to you.

Start with the quick wins, where impact is high and effort is low. Let their results fund the strategic bets once once your data is ready.

HOW PRAXENT HELPS

Praxent's AI-native bank matrix scores 72 banking workflows across eight areas by effort and P&L impact, so your shortlist starts from a scored view.

Open the interactive matrix →

Read the blog post →

PILLAR 6

The 30 / 90 / 365 roadmap

Scale from first use cases to bank-wide results, one proven phase at a time



Only 37% of AI initiatives launched in the past year are live and delivering value.[6]

Which means nearly two-thirds of them failed to show their worth. One way to avoid that fate is by tying each AI dollar to proof, and starting the next phase only when the last one has been delivered. That way your executive sponsor can show the board the returns before committing more budget.



37%

of AI initiatives launched in the past year are live and delivering value

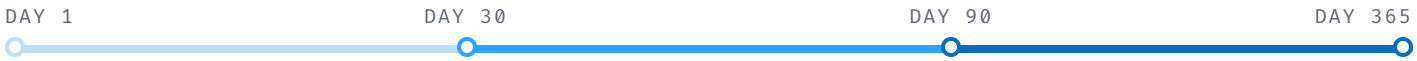
S&P Global, June 2026

Your bank moves to the next phase only after the last one shows results

Set target dates for each phase, and move only once the evidence is in. Each new wave of teams starts with people from the last one, so the next group hears what worked from colleagues they trust.

THE 30 / 90 / 365 ROADMAP

You leave the Foundation phase with weekly adoption you can measure, named champions, and a scored list of the workflows worth taking on next.



Foundation DAYS 1 TO 30

- ✓ AI manifesto and executive team
- ✓ First use cases scored
- ✓ Champions launched

◇ Gate 1 · Foundation complete
Measurable weekly adoption, named champions, and a scored backlog

Adoption DAYS 31 TO 90

- ✓ 80% weekly active users
- ✓ First use cases deployed
- ✓ Workflow audit complete

◇ Gate 2 · Scale begins
The next wave opens only when this one proves value

Scale DAYS 91 TO 365

- ✓ 80% weekly active, proficient
- ✓ One agent per employee
- ✓ First board AI report

Each wave is seeded with graduates who open with their own wins.

By day 365, your own team runs AI across the bank

Record your baseline numbers before day one, so your day-365 report to the board measures a full year against where the bank started.

Once your team runs its flagship workflows on its own, open access across the bank with a center of excellence and a certified trainer network.

FIVE MOVES TO CARRY AI ACROSS YOUR BANK

01	Redesign your flagship workflows inside your own systems	Work through each one with your partner, and confirm your team can run it alone before the partner steps back.
02	Roll out what works to every business line	Each line starts from workflows another part of the bank has already proven.
03	Launch a center of excellence	A small team keeps the scored list current and moves proven workflows between business lines.
04	Certify in-house trainers	Your own trainers keep the program running long after outside help ends.
05	Write a graduation brief for every agent that moves into daily use	Each brief records the procedure the agent follows, the return it has measured, and when it hands a case to a person.

⌘ HOW PRAXENT HELPS

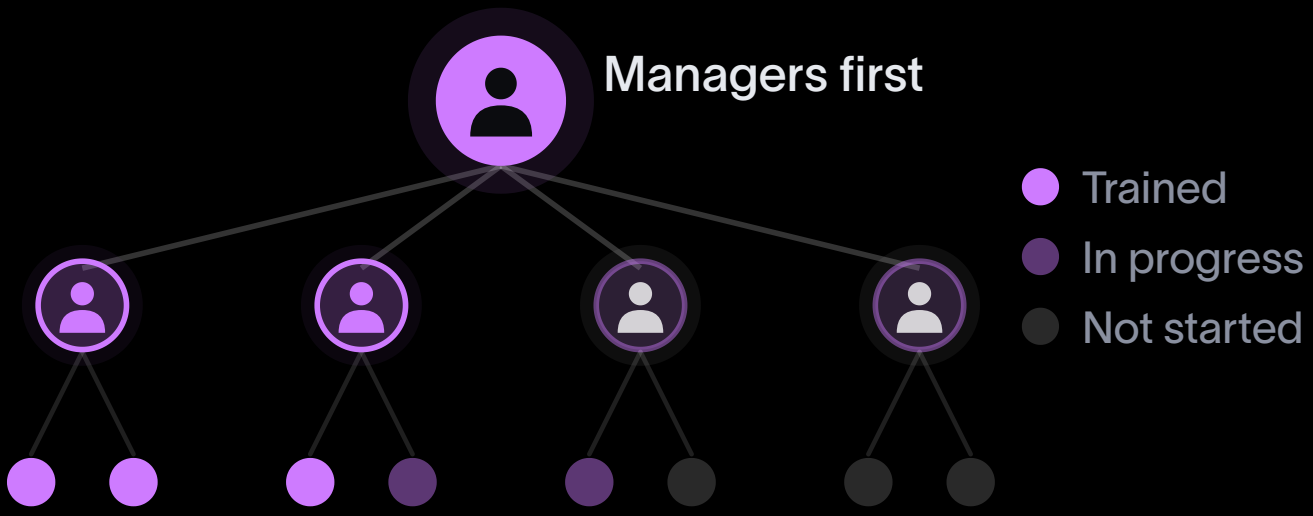
Praxent runs the roadmap with your team, starting with a first group using AI on their own work and measuring gains against the baseline you set before day one. Your sponsor gets the evidence to open each phase.

Plan your first year →

PILLAR 7

Upskilling

Train your managers first, because their expectations set your return on AI



Once your people have AI tools, their skill level is what decides your ROI.

Middle managers (38%) and senior leaders (35%) are the least AI-ready groups in their organizations,[7] yet they set the expectations your AI spend depends on.

50%

Higher worker proficiency scores when their managers expect and require AI use.[13]

AI PROFICIENCY TODAY

33%

US workers feeling overwhelmed by AI

61%

Use AI only for search[13]

5.5%

US workers who are AI proficient[13]

2.3/10

Average prompting score

Start by training managers who set the pace first.
Then make AI use part of every role's review.

This three-course ladder moves every role up one proficiency level at a time.

Each course meets people where they are as 101 takes a novice to experimenter, 201 takes an experimenter to proficient, and 301 takes a proficient user to expert.

PROFICIENCY LADDER

LEVEL 1

STARTING POINT

Novice

Knows AI exists but hasn't used it in a meaningful way at work

LEVEL 2

COURSE 101

Experimenter

Uses AI now and then for basic tasks such as summarizing, drafting, and search

LEVEL 3

COURSE 201

Proficient

Builds AI into daily work and writes effective prompts for complex tasks

LEVEL 4

COURSE 301

Expert

Builds custom assistants, automates workflows, trains others, and finds new uses

Four rules every employee learns before touching customer data

Before anyone works with customer data, training also covers the banking red lines that informal, ad hoc learning won't reach. Every course is built on the work your people already do, so the gains show up in their own tasks. Business teams learn on Claude Cowork, and engineering teams learn on Claude Code.

Responsible AI in regulated work

What they can and can't do with customer data

AI disclosure

When customers must be told AI was involved

Compliance red lines

The work that AI must never produce on its own, such as BSA filings and adverse action notices

Secure use

Customer personal information never goes into unapproved AI tools

HOW PRAXENT HELPS

Praxent's Claude Cowork and Claude Code enablement programs take your people through the 101, 201, and 301 courses on their own work, with banking rules built in.

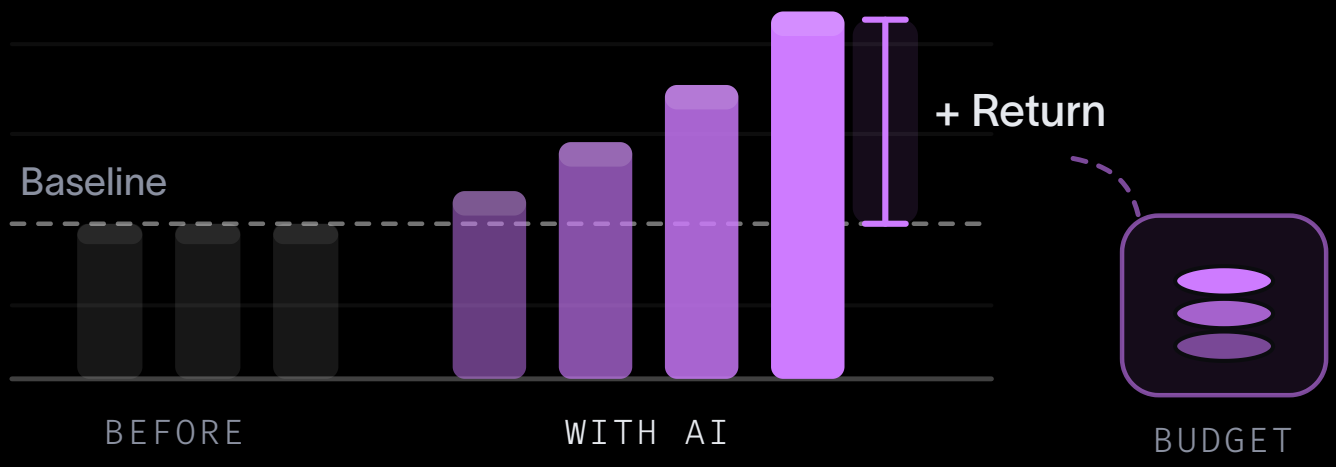
Start with a one-day Cowork Hackathon, where your team hands one of its own procedures to AI and leaves with a ranked list of the next tasks worth moving.

Plan your training →

PILLAR 8

Measurement and budget

A baseline that proves AI's return so you can budget for it



Only 26% of engineering leaders say their organization has a robust way to measure AI's business value.[14]

Prove AI's return to your board using the benchmarks they already track.

Record your baseline numbers before AI goes live, and measure each gain against them. Use cases that can prove their value earn a budget line so you can keep funding what works.



26%

of engineering leaders say their organization has a robust way to measure the business value of AI[14]



56%

called their institution's AI budgeting "guesswork"[14]

Measure personal wins in months 1 to 6, with workflow wins in months 3 to 12

Personal wins, such as team sentiment, output per employee, and productivity, are what keep people using AI at the start. This allows time for the workflow results that fund the program to take shape.

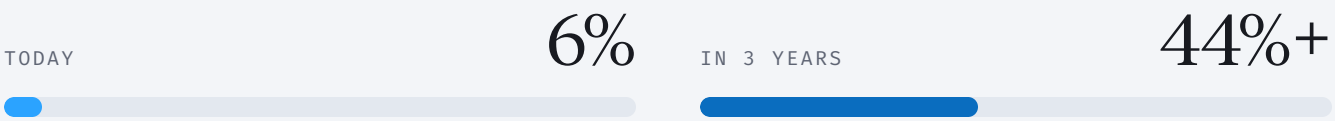
Starting in month 3, track workflow wins against your baseline and turn each 90-day sprint's gains into a value story your board can follow.

Track speed and quality together on five core bank workflows

Pair each speed metric with its quality check, so speed never hides a quality problem. Use these pairings as a starting point for your value stories.

SPEED METRIC	QUALITY CHECK TO PAIR IT WITH	BUSINESS RESULT IT SUPPORTS
Loan processing time	Exceptions and rework per file	Faster funding and more loans per underwriter
BSA alert disposition time	Closed alerts that quality review or independent testing reopens, and SARs filed on time	Investigators spend their hours on alerts that need a filing decision
Customer inquiry resolution time	Repeat contacts and complaints	Lower cost-to-serve and more satisfied customers
Compliance review cycle time	Issues found later by internal audit or examiners	Faster launch of new products and campaigns
Report generation hours	Corrections after a report goes out	Staff time moved from assembling reports to analysis

AI'S SHARE OF BANK TECHNOLOGY BUDGETS



Financial institutions expect AI to take up to 6% of current technology budgets, and most believe that share will grow past 44% over the next three years.[3] Your 2027 number becomes the floor for every year after it.

Budget for five line items, each adjusted for what your bank needs

Most of your AI budget lines will carry compliance and risk costs like model validation and security tooling that a typical AI budget leaves out. We compare the typical approach and what changes for a bank below. This way, you can plan for the costs up front and keep the best use cases from stalling in review.

Here's what that looks like in banking.

LINE ITEM	TYPICAL APPROACH	FOR BANKING
AI platform licenses	Per-person licenses, such as Microsoft 365 Copilot at $30 per user per month at enterprise list price	A higher tier for security and compliance features
Training and enablement	A dedicated program	Compliance-specific AI training
Head of AI and team	A dedicated function	Compliance and risk members on the team
Workflow automation sprints	3 to 5 a year	Model validation for each use case
Infrastructure and integration	Platform and connection costs	Data integration, security tooling, and on-premises options

HOW PRAXENT HELPS

Praxent offers outcome-based engagements tied to the results in your value stories, such as minutes saved per task, rework rate, and cycle time, with every gain reported against your baseline.

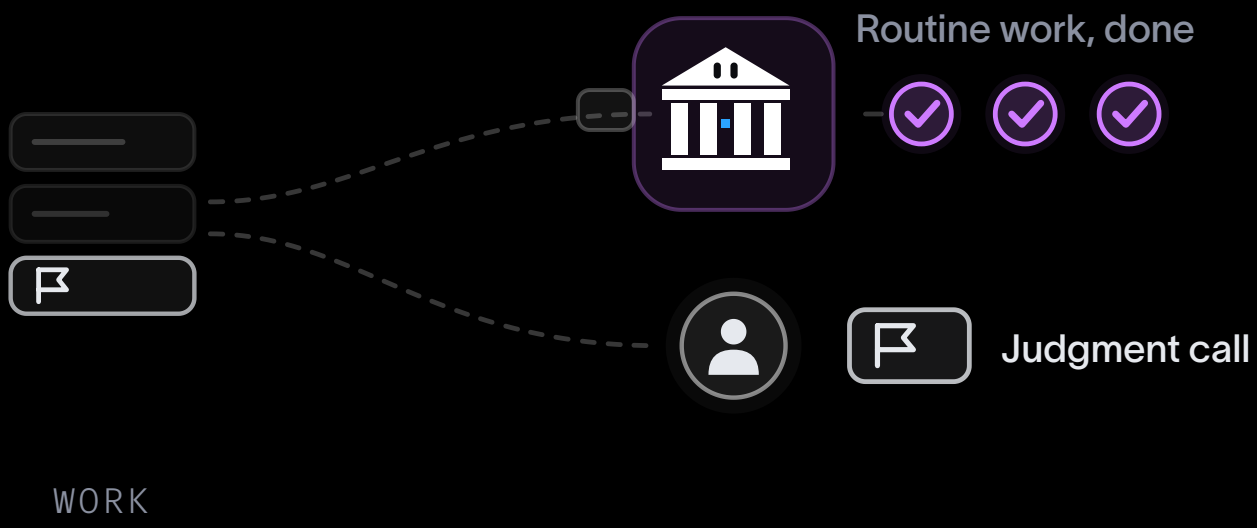
Your initiative may also qualify for Microsoft and Praxent partner funding that lowers what your bank spends to start.

Build your AI budget →

PILLAR 9

Agent strategy

Hand routine work to agents so your people can focus on the judgment calls that grow revenue



When agents take on organizing and checking each loan file's documents, your underwriters can spend their day on credit decisions. And by having agents clear the overnight alert queue, your BSA investigators can start their morning on the few cases that need review.

AI-native banks run on leaders, managers, and employees working alongside purpose-built agents.

Plan for one agent per employee by the end of year one.

Only 4% of banks have agents fully integrated across their organization,[3] so a governed agent program puts your bank among the first to run them at scale.



4%

of banks have agents fully integrated across their organization

S&P Global Ratings, 2026

Agents can act on their own only after passing three stages

STAGE 1

Supervised automation

WHAT IT LOOKS LIKE

The first workflows run with a person supervising

WHAT IT TAKES

- ✓ Four safety tests passed before deployment

STAGE 2

Managed fleet

WHAT IT LOOKS LIKE

A growing set of agents runs routine work, each one documented and owned

WHAT IT TAKES

- ✓ A written procedure
- ✓ A named owner
- ✓ An off-switch you've tested for every routine

STAGE 3

Production agents

WHAT IT LOOKS LIKE

Agents carry daily work and hand exceptions to people

WHAT IT TAKES

- ✓ A stable procedure
- ✓ A measured return
- ✓ A defined path to a person, recorded in the agent's graduation brief from Pillar 6

Before any agent goes live, decide what it can access and confirm that every action it takes can be undone. An agent must prove itself at each stage before the next can go.

4 safety tests each agent must pass before going live

Govern agents the way you govern models, with an approved inventory, permission tiers that follow your data classifications, an audit trail for every action, and human review scaled to each agent's risk.

Predetermined security

WHAT IT REQUIRES

What the agent can access is set before it runs

WHAT IT PROTECTS

Customer data stays inside the boundaries you drew in Pillar 3

Reversible

WHAT IT REQUIRES

Every action the agent takes can be undone

WHAT IT PROTECTS

A wrong action can be rolled back in full

Portable

WHAT IT REQUIRES

The workflow can move to another platform

WHAT IT PROTECTS

Your bank isn't locked into one vendor

Measurable

WHAT IT REQUIRES

Results are tracked against your baseline

WHAT IT PROTECTS

Your board can see what each agent returns

These 5 controls keep a person in the loop where risk is highest

01	Keep an approved inventory of every agent.	Each one has an owner, a purpose, and a record of what it's allowed to do.
02	Set permission tiers by data access.	An agent's access follows the data classifications you set in Pillar 3.
03	Monitor every agent and keep an audit trail.	You can trace every action an agent took back to the request behind it.
04	Scale human review to each agent's risk.	Agents that touch lending decisions, BSA work, or customer communications keep a person in the loop.
05	Review performance and accuracy on a set schedule.	Pause any agent whose accuracy falls below its baseline until it's fixed.

With these controls, you'll put every agent on record. Keep these records and each agent's off-switch test in your Pillar 3 show file so it's ready for any questions from regulators.

HOW PRAXENT HELPS

The Agent Proof of Concept shows your board one of your workflows run by an agent, with its accuracy, speed, and cost measured against cases your team has already decided.

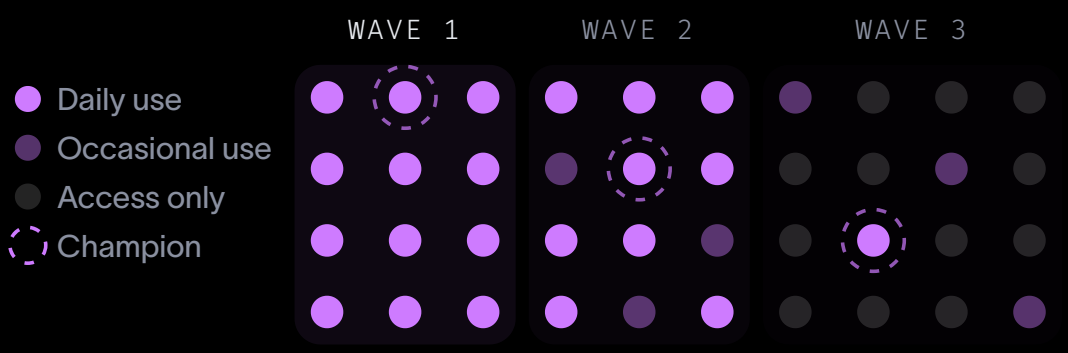
You also get a validation package your model risk team can sign off on before any agent leaves Stage 1.

Pick your first agents →

PILLAR 10

Change management

Turn your team's AI access into daily use



Your best return on AI depends on whether your people choose to use it every day. Teams adopt faster when they know why it's coming and what it means for their role. Unfortunately, only 10% of executives say their organization explains that clearly.

AI champions change the culture starting with their own teams

You need internal advocates that spur team members forward. Department leads nominate should pick champions who are respected by their peers, curious, carrying a full workload, and willing to teach. With Praxent's Claude Cowork enablement, Champions earn the right to co-teach by passing a scored teach-back, and each cohort gets a charter, an office-hours rotation, and a shared channel.

Begin where there's a willing team, respected voices and usable data. Open the next wave only once the current one hits its adoption target and files its scorecards.

HOW PRAXENT HELPS

Praxent's Claude Cowork enablement rolls out team by team with compliance checks built in. We stay until your champions run the program and your board sees the return against your baseline.

Turn access into use →

This 10 pillar plan closes the main gaps blocking regular use

Six moves to make in your first 30 days to become a market-leading bank

01 Draft the AI manifesto and put it on the next board agenda.

02 Name your five AI owners and form the governance committee.

03 Score 3 to 5 use cases on effort and P&L impact.

04 Launch the champions program in your first wave of departments.

05 Run the first enablement sessions, with every attendee working on a task from their own job.

06 Assess the regulatory risk of your first use cases with risk and compliance in the room.

Your first month sets up every target below.

Five targets for a market-leading, AI-native bank at one year

These are the 12-month targets Praxent sets with banks becoming AI native, so your board sees progress at 90 days and at a year.

12-MONTH TARGET

Measured gains against your baseline	80%	80%	5%	1:1
	of your workforce	of your workforce	of the workforce	agent per employee
MEASURE	MEASURE	MEASURE	MEASURE	MEASURE
Cycle time and cost-to-serve on your first three workflows	Weekly active AI users	Workforce at proficient (Level 3) or above	AI champions across functions	Agent-to-employee ratio

YOUR BANK'S BEST STARTING POINT

Start with the result your board wants first, whether that's a modern core, a faster back office, or a more personal customer *experience*

AI earns its budget one proven result at a time, so pick the pressure point your board raises most and show a business result before you commit to scale.

01

IF YOUR BOARD ASKS ABOUT

Modernizing the core without losing its rules

START WITH

PraxentDecode℠

WHAT YOU'LL BE ABLE TO SHOW

Exactly what your core does, so you can modernize around it and put AI to work on it without losing the rules it holds

2 TO 4 WEEKS

02

IF YOUR BOARD ASKS ABOUT

Manual back-office queues

START WITH

Agent Proof of Concept

WHAT YOU'LL BE ABLE TO SHOW

One workflow run by an agent, with its accuracy, speed, and cost against your team's past cases and a validation package compliance can sign off on

4 TO 8 WEEKS

03

IF YOUR BOARD ASKS ABOUT

Deposit growth and the customer experience

START WITH

CX Prototype

WHAT YOU'LL BE ABLE TO SHOW

A customer-tested design for a digital experience that turns more applicants into funded accounts, before you commit to the build

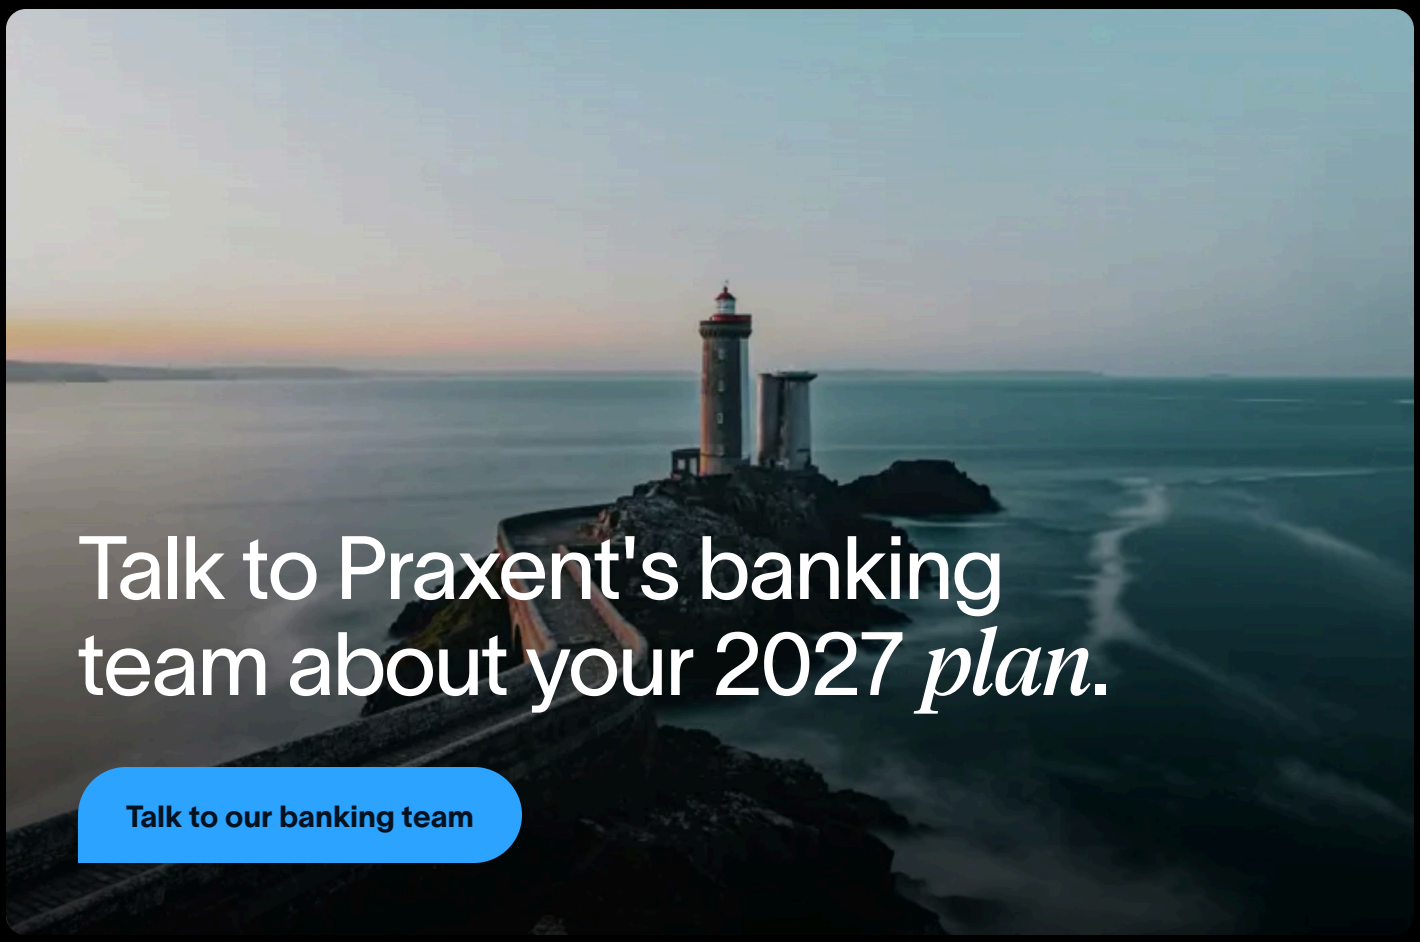
6 TO 8 WEEKS

Every AI system Praxent delivers is built for examiner sign-off. It runs on PraxentGuard℠, our AI control framework, with full audit trails, SOC 2 Type II alignment, and human review on every decision that matters.

Praxent's banking expertise has impressed me the most. I did my due diligence and searched for vendors with deep banking experience. Praxent had this.

Aldo Pietropaolo

Chief Solutions Officer, Neocova



Praxent is an award-winning AI transformation firm for financial services and a Select Services Partner in Anthropic's Claude Partner Network. With more than 25 years in financial services and 200+ fintech AI experts, we help banks and fintechs become AI-native market leaders.

Praxent

Select Services Partner
Claude Partner Network

Microsoft Solutions Partner
Cloud & AI Platforms

Endnotes

1 Miriam Fernández and Nicolas Charnay, "AI and banking: Leaders will soon pull away from the pack," S&P Global, October 28, 2025

2 S&P Global Market Intelligence 451 Research, "Voice of the Enterprise: AI & Machine Learning, Use Cases 2025," January 2025, as cited in note 1. Financial services firms abandon an average of 48% of AI projects before production

3 S&P Global Ratings, "AI In Financial Institutions Survey: Uneven Adoption And Risk Management Will Drive Credit Quality Differentiation," September 21, 2026. Survey of 179 financial institutions

4 Board of Governors of the Federal Reserve System, SR 26-2, "Revised Guidance on Model Risk Management", April 17, 2026. Supersedes SR 11-7.

5 Office of the Comptroller of the Currency, News Release 2026-29 and Bulletin 2026-13, "Model Risk Management: Revised Guidance", April 17, 2026. Rescinds OCC Bulletin 2011-12.

6 Pollyanna De Lima, Alexander Johnston, and Sophie Malin, "The AI and labor landscape 2026," S&P Global, June 2, 2026, drawing on 451 Research survey data

7 HFS Research with EY, "The operating model tipping point," September 10, 2026. Survey of 302 senior executives across the Global 2000

8 Interagency Guidance on Third-Party Relationships: Risk Management, OCC Bulletin 2023-17, June 2023. A replacement was proposed September 11, 2026, with comments due November 16, 2026.

9 CFPB, Regulation B final rule, April 22, 2026, effective July 21, 2026

10 Lending law, disparate impact standards

Fair Housing Act

U.S. Supreme Court, Texas Department of Housing and Community Affairs v. Inclusive Communities Project, Inc., 576 U.S. 519 (2015). The Court held that "disparate-impact claims are cognizable under the Fair Housing Act."

42 U.S.C. § 3605, "Discrimination in residential real estate-related transactions." This section is what makes the Fair Housing Act a banking issue. It covers "the making or purchasing of loans" to buy, build, or repair a dwelling, and loans "secured by residential real estate." Meaning it reaches mortgage and home equity lending.

U.S. Department of Housing and Urban Development, "HUD's Implementation of the Fair Housing Act's Disparate Impact Standard," proposed rule, 91 FR 1475, January 14, 2026, and the "supplemental proposed rule," 91 FR 51416, August 10, 2026.

New York

New York State Legislature, Assembly Bill A4040-A / Senate Bill S4067-A, "Codifies the disparate impact standard in the human rights law," signed December 19, 2025, as Chapter 649 of the Laws of 2025.

11 Illinois SB 3777, effective January 1, 2027
Alan S. Kaplinsky, Richard J. Andreano, Jr., and John L. Culhane, Jr., "Illinois Enacts Disparate Impact Standard for Credit Decisions: What Lenders Need to Know," Consumer Finance Monitor (Ballard Spahr), August 14, 2026.

12 Nupur Anand, Phoebe Seers and Tommy Reggiori Wilkes "Banking watchdogs ramp up AI scrutiny as stability fears mount, sources say" Reuters, October 13, 2026, on the Fed, OCC, and FDIC asking banks about their use of AI agents

13 Section, AI Proficiency Report, June 2026, and Section, January 2026 press release (Business Wire)

14 Harness, "State of AI in FinOps Report 2026," survey of 700 engineering leaders and practitioners conducted in May and June 2026.

15 Pew Research Center, "U.S. Workers Are More Worried Than Hopeful About Future AI Use in the Workplace"

16 Thomson Reuters, "Future of Professionals Report 2026," June 22, 2026. Survey of 1,816 professionals across law, tax, audit, accounting, compliance, risk, and global trade

17 Section blog, "85% of workers don't use AI for business value"