Architecture Overview

Agents sit on top of your existing case files. They keep each case's workflow current, draft routine documents, catch errors, and propose changes for a person to approve. One runtime powers every engine you deploy, so you start with one workflow and add more over time.

YOUR TEAM
People stay in charge
Work cases and queues, review drafts, and approve or reject every proposed action. Every result shows why.

AGENTS DASHBOARD
One place to manage agents
Live pipeline by stage, approvals queue and a full activity record across every engine.

YOUR IDENTITY PROVIDER
Your sign-in, your roles
Company sign-in with roles and permissions from the groups you already manage.

↓ RUNS ON

Your deployment

CONFIGURED TO YOUR INSTITUTION

Your workflows, roles, approval policies, templates and branding, configured on the engines you choose.

Your AgentCommand℠



Run several engines side by side, all managed from one agents dashboard: every pipeline, approval and agent action in one view. The agents are yours.

- Workflows fitted to how your team works
- Approval policies per action type
- Templates in your format and branding
- Connections to your systems

↓ BUILT FROM

Engines

40+ ENGINES · SIX INDUSTRIES

One engine per workflow, each carrying the domain knowledge for that process. Deploy one, then add more as you grow.

LENDING 7
- Commercial Lending
- Mortgage Lending
- Auto Lending

BANKING 7
- Consumer Account Opening
- BSA/AML
- Treasury Operations

PAYMENTS 5
- Merchant Onboarding and Underwriting
- Chargeback Operations
- Settlement and Reconciliation

INSURANCE 5
- Commercial Lines Underwriting
- Commercial Claims
- Life and Annuity New Business

HEALTH AND BENEFITS 9
- PBM Pricing
- Prior Authorization and Appeals
- Claims

INVESTING 9
- Fund Operations
- PE Due Diligence
- LP Reporting

YOUR NEXT ENGINE

Need a workflow that isn't in the library? We build it as a new engine on the same runtime, fitted to your process.

↓ POWERED BY

Core runtime

SHARED BY EVERY ENGINE

01
Connect
Reads documents and case status from your systems.

02
Understand
Classifies documents and extracts facts, with confidence.

03
Decide
Your rules set each step's status. Models never decide.

04
Act with approval
Agents draft and propose. Your people approve.

05
Record
Every step and approval is logged for audit.

↓ INSIDE YOUR ENVIRONMENT

YOUR SYSTEMS
Your files and systems of record
Core, loan and policy systems, CRM, documents and inbox. Changed only through approved actions.

YOUR AI PLATFORM
Azure AI Foundry or the Anthropic API
Claude or another approved model, pinned to the version your risk team signed off.

YOUR CLOUD
Azure, AWS or Google Cloud
Runs in your environment, with your data in your database. Azure-native, built to pair with Foundry.

Any model. Any cloud. Your *controls.*

AgentCommand℠ runs on the model and cloud your institution already trusts. The governance and audit trail stay the same wherever you deploy.

MODELS AND AI PLATFORMS

Built for Claude on Azure AI Foundry.

AZURE AI FOUNDRY

Claude inside your Azure tenant

Run Claude and other Foundry models under the Azure controls, security and agreements you already have. Each deployment pins the exact Foundry model your risk team approved.

ANTHROPIC API

Claude, direct

Anthropic's frontier models for extraction, analysis and drafting, called straight from AgentCommand℠.

ALSO WORKS WITH AMAZON BEDROCK GOOGLE VERTEX AI Same governance on every platform.

CLOUDS

Deployed in your cloud.

MICROSOFT AZURE

Azure-native from day one

Azure AI Foundry for models, Entra ID sign-in, Key Vault for secrets, private endpoints, Azure SQL and Application Insights monitoring.

AWS AND GOOGLE CLOUD

The same runtime, your platform

Deploy on the cloud you run, with your own storage, secrets and identity provider.

CONTROLS

Ready for your review board.

You own your agents

Your agents, your data and every improvement they learn stay with your institution.

People approve every write

Agents propose and never write directly. A policy per action type sets who approves.

Separation of duties

Nobody approves their own work, and low-confidence results always go to a person.

A record of everything

Every agent step and approval is logged with its sources, so you can answer any audit question.

Private by default

Prompts and documents stay in your database. Monitoring carries metadata only.

Examiner ready

Model-risk documentation is generated for your SR 26-2 review.

Praxent

Specialists in Claude and Azure for financial *services.*

Praxent is Anthropic's Select partner dedicated to financial services and a Microsoft Solutions Partner for Cloud and AI Platforms. Our forward-deployed fintech engineers map your process, launch each engine with your compliance team, and keep tuning it to drive ROI.

Select Services Partner
Claude Partner Network

Microsoft Solutions Partner
Cloud & AI Platforms

25+
years in financial services

100+
Claude Certified Architects

SOC 2
Type II

Book a demo hello@praxent.ai