The Financial Services Guide to Building Autonomous AI Agents

ANTHROP\C | aws

Contents

Executive summary

Your team is stuck between two worlds.

Institutions face mounting pressure from aging infrastructure requiring costly modernization, digital-native competitors capturing market share, and customers expecting real-time, personalized financial experiences. Once hallmarks of prudent risk management, manual technologies and processes now create operational bottlenecks that limit growth and innovation.

Here's the good news: teams using AI agents are already transforming these pain points into competitive advantages. Financial services leaders like Verisk and NBIM are helping millions of customers at once, catching fraud before it happens, and turning weeks-long loan processes into same-day approvals.

In this guide, we'll demonstrate how to build secure, scalable AI agents using Claude, Anthropic's family of state-of-the-art foundational models designed for safety and reliability, deployed through Amazon Bedrock's fully managed generative AI infrastructure.

What you'll learn:

- **How financial services organizations use autonomous AI agents** for fraud detection, customer service, and more
- **Step-by-step blueprints** for building agents with Claude in Amazon Bedrock that work within your existing systems
- **Realistic timelines and team structures** based on what actually works
- **Stories from Verisk and NBIM** showing real results, not just promises

Ready to move beyond AI experiments to production-scale deployments that meet the compliance, security, and performance standards your financial institution requires?

Let's dive in.

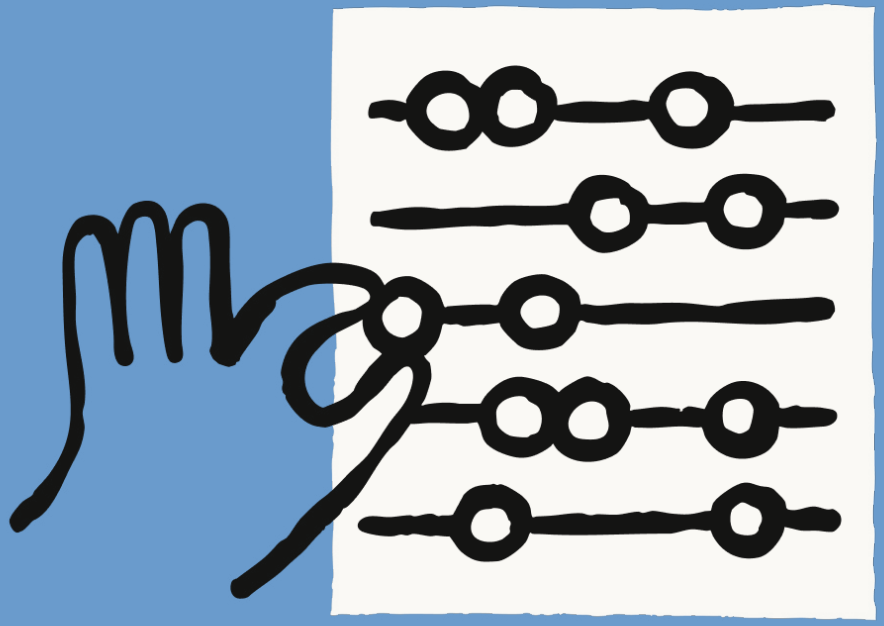

Chapter 1

Why financial services institutions need AI agents (now)

Chapter 1

Why financial services institutions need AI agents (now)

Most financial services teams are drowning in manual work while trying to meet rising customer expectations, and existing solutions aren't fitting the bill. The gap between what you need to deliver and what your current tools allow keeps growing—and it's costing more than just money.

The cost of the status quo

Manual approaches to financial services operations carry hidden costs that compound annually:

Your people are buried in repetitive tasks

Think about your last commercial loan application. How many systems did it touch? How many people had to manually move data between them? Each handoff adds hours, introduces errors, and frustrates both your team and customers. When you multiply this across thousands of transactions, you're looking at millions in preventable costs—and burned-out employees who could be doing more valuable work.

Customers leave for better experiences

Your customers can split a dinner bill instantly on their phones, yet opening a new account with you takes days. They're not being unreasonable—they're comparing you to every other digital experience they have. Each friction point pushes them toward competitors who've figured out how to make finance feel effortless. And with customer acquisition costs through the roof, losing loyal customers because of clunky processes hits twice as hard.

Compliance keeps getting more complex (and expensive)

Regulatory requirements grow more intricate every year, and manual compliance processes can't keep up. Beyond the direct costs of fines, your compliance teams spend most of their time on repetitive checks instead of catching sophisticated threats. As regulations evolve, throwing more people at the problem becomes unsustainable.

How autonomous AI agents change the game

Instead of patching each problem with point solutions, leading teams are deploying AI agents that tackle interconnected challenges while working within existing constraints.

The key? Start small. Pick one specific pain point, prove the value, then expand. The teams seeing real results aren't trying to transform everything overnight—they're building confidence with targeted wins.

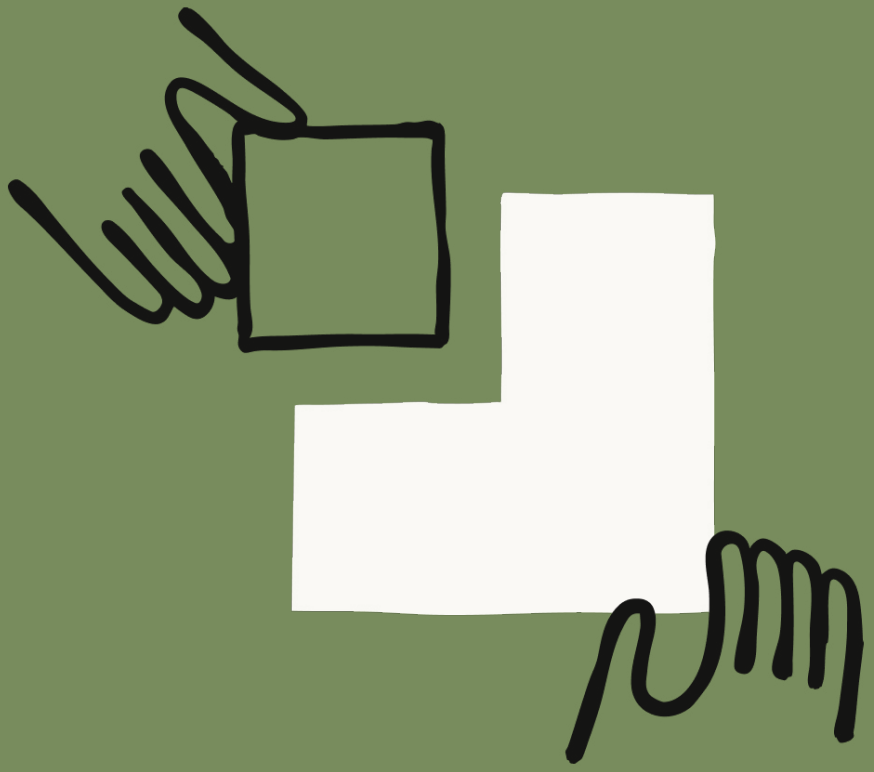

Chapter 2

Better together: How Anthropic and AWS power autonomous AI agents

Chapter 2

Better together: How Anthropic and AWS power autonomous AI agents

Building AI agents for financial services requires exceptional precision and accountability. Your agents need to understand complex regulations, handle sensitive data securely, and explain every decision they make. One wrong calculation or data breach can trigger serious consequences.

This is why leading financial institutions are turning to **Claude**, Anthropic's family of advanced AI models known for their industry-leading coding and reasoning capabilities, nuanced understanding, and unwavering emphasis on safety and reliability. Claude doesn't just process data; it understands context, maintains consistency across complex workflows, and provides the explainability that regulators demand.

Deployed through Amazon Bedrock, AWS's fully managed service that delivers enterprise grade security, privacy and responsible AI capabilities, Claude becomes even more powerful. Bedrock eliminates the infrastructure complexity that typically slows innovation, while providing the robust governance controls and compliance frameworks that financial services require.

In this section, we'll share core capabilities of Claude and Amazon Bedrock, as well as additional services and solutions that support agentic development on AWS, including:

- **Anthropic's Financial Analysis Solution (FAS)** for building custom financial services agents
- **Model Context Protocol (MCP)** for providing agents with secure access to relevant data and tools
- **Strands Agents SDK** for building and running agents with minimal code
- **Amazon Bedrock AgentCore** IaaS for deploying and operating agents at scale

Together, Claude in Amazon Bedrock represents more than just technology. It's a complete solution that enables financial institutions to build powerful, safe agents at scale without compromising on the security, compliance, and reliability that define trust in financial services.

Choose the right Claude model for your needs

The Claude family offers a range of models to match different financial services use cases and performance requirements:

Claude Haiku models provide lightning-fast responses for high-volume, cost-sensitive applications like customer support and document processing.

Claude Sonnet models balance strong performance with efficiency, making it ideal for complex tasks like financial analysis and risk assessment.

Claude Opus models deliver maximum intelligence for the most demanding applications requiring deep reasoning and comprehensive analysis.

At the core of Claude's design is a commitment to transparency and trustworthiness, critical requirements for financial services where decisions must be auditable, compliant, and aligned with regulatory standards.

Building trust with explainable AI

Your stakeholders need to trust AI decisions. Claude's **Constitutional AI framework** ensures every decision can be explained and audited. With SOC 2 Type II certification and a guarantee that your data never trains the model, you can deploy AI without compromising security or intellectual property.

Achieve superior financial reasoning and analysis

Claude's reasoning capabilities set it apart in the complex world of financial analysis, where nuanced understanding can mean the difference between profit and loss. Claude Opus 4.1 demonstrates **top performance** on the Finance Agent Benchmark, outperforming competitors on critical financial reasoning tasks, from analyzing earnings transcripts to building complex valuation models.

Norges Bank Investment Management

Agents in the wild: The Norwegian sovereign wealth fund experienced 20% productivity gains (equivalent to 213,000 hours saved) with Claude in Bedrock. **Learn more here**.

Deploy Claude at enterprise-scale with Amazon Bedrock

For financial institutions seeking to harness Claude's capabilities while maintaining enterprise-grade security, performance, and integration with existing infrastructure, Bedrock emerges as the purpose-built platform that bridges cutting-edge AI with the stringent demands of financial services.

Ensure security and meet compliance requirements

Bedrock keeps your data within your secure AWS environment. Your information never leaves your control or gets used to train models. With private endpoints in your VPC, you can connect Claude to internal systems while maintaining complete isolation—critical for sensitive use cases like anti-money laundering analysis.

Agents in the wild: With Claude in Bedrock, Brex automated 75% of expense transactions with a 94% compliance rate, vs. the industry standard of 70%. Learn more here.

Scale seamlessly and optimize performance

Amazon Bedrock also transforms how financial institutions deploy and manage AI models like Claude at scale, eliminating the complexity of infrastructure management while delivering consistent, high-performance AI inference.

Agents in the wild: IG Group experienced triple-digit improvements in speed-to-market for their global trading platform and financial services products with Claude. Learn more here.

Accelerate financial analysis with Anthropic's Financial Analysis Solution

In addition to using Claude in Bedrock, financial institutions can leverage Anthropic's **Financial Analysis Solution** (FAS) as a production-ready foundation for building AI agents specifically geared towards financial services use cases. FAS accelerates agent development through pre-built integrations with over 10 critical financial data providers:

- **Market intelligence:** Direct connections to FactSet, S&P Capital IQ, and other data sets enable agents to access real-time market data, corporate fundamentals, and historical analytics without custom API development.
- **Deal and company data:** Integrations with Daloopa, Morningstar, and PitchBook allow agents to automatically pull deal comparables, financial metrics, and industry benchmarks for due diligence workflows.
- **Enterprise analytics:** Native connections to Databricks, Snowflake, and Box ensure agents can seamlessly access proprietary data lakes, financial models, and internal research.

Beyond connectivity, the solution equips financial agents with specialized tools:

- **Native Office Integration:** Agents can directly manipulate Excel models and generate PowerPoint presentations, eliminating manual data transfer between AI insights and deliverables
- **Agentic Research Workflows:** Pre-configured patterns for equity research, credit analysis, and portfolio optimization that combine web search, document analysis, and quantitative modeling
- **Financial Prompt Library:** Battle-tested prompts for DCF modeling, comparable analysis, earnings interpretation, and risk assessment that encode best practices from leading financial institutions

Provide contextual awareness with Model Context Protocol

Another technology financial services organizations can use to build autonomous AI agents is Model Context Protocol (MCP), an open framework developed by Anthropic that enables financial agents to connect and interact with external tools, applications, and data sources.

MCP's true power emerges when orchestrating multiple specialized agents into collaborative intelligence networks. Financial institutions can create agent teams where each member brings specific expertise.

Consider a private equity due diligence workflow:

1. **Market Analysis Agent** connects to a FactSet MCP server to gather sector trends and comparable transactions
2. **Financial Modeling Agent** accesses internal data warehouses through enterprise MCP servers to build valuation models
3. **Risk Assessment Agent** queries compliance databases and regulatory feeds to identify potential issues
4. **Synthesis Agent** aggregates insights from all specialists to generate investment committee materials

Build flexible agentic systems with Strands Agents

Strands Agents is an open source SDK that accelerates how financial teams can build Claude-powered autonomous AI agents by eliminating boilerplate code and focusing on business logic. Its model-driven architecture particularly suits financial analysis where requirements evolve rapidly with market conditions.

This simplicity enables financial teams to iterate rapidly, adjusting agent behavior for new regulations, market conditions, or investment strategies without rebuilding infrastructure.

Deploy agents in production with Amazon Bedrock AgentCore

Amazon Bedrock AgentCore provides purpose-built services and tools to handle dynamic agent workloads, offering capabilities like secure identity management, memory, code interpretation, web automation, gateway and observability, all working together to eliminate the need to manage complex agent infrastructure.

While complete session isolation ensures that agents handling different clients' data never intersect, integration with existing identity providers enables automated permission management aligned with institutional access controls.

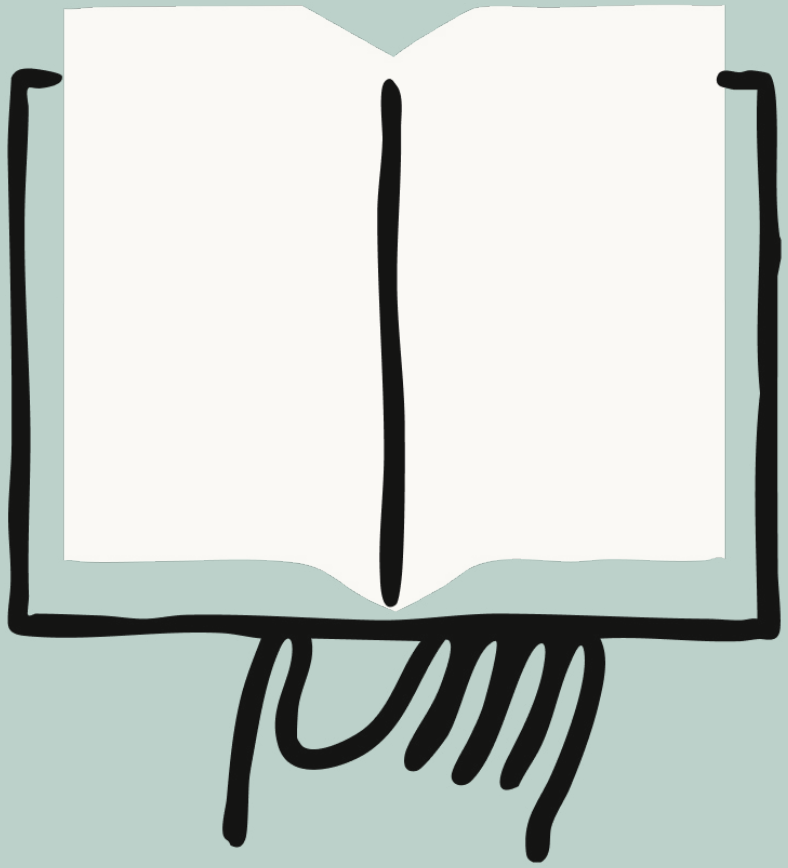

Chapter 3

Agentic use cases for financial services

Chapter 3

Agentic use cases for financial services

With the right agentic foundations in place, financial services organizations can take advantage of autonomous AI agents across a variety of high-impact use cases, including fraud detection and prevention, customer service automation, portfolio optimization and advisory and code modernization.

Fraud detection and prevention

Your fraud team faces an impossible task: analyzing millions of transactions in real-time while fraudsters constantly evolve their tactics. AI agents change the equation.

Real-time transaction monitoring leverages Claude in Bedrock to process millions of transactions per second with sub-100ms response times. The architecture combines Bedrock's streaming inference with Amazon Kinesis Data Streams to analyze transaction patterns, merchant profiles, and customer behavior instantly. AWS Lambda triggers Claude's analysis on high-risk transactions while Amazon DynamoDB maintains real-time fraud scores, enabling instant decisions that balance security with customer experience.

Pattern recognition and anomaly detection extends beyond traditional rule-based systems through Claude's multi-dimensional analysis capabilities. Using Bedrock's batch inference, Claude analyzes historical data in Amazon Simple Storage Service (Amazon S3) to identify emerging fraud patterns, then deploys these insights through real-time monitoring agents. The system incorporates unstructured data—i.e. customer communications, social media signals and dark web intelligence—creating an adaptive fraud detection network.

Customer service automation

Your service team knows every customer interaction shapes loyalty. AI agents help deliver the instant, personalized service customers expect.

Intelligent routing and resolution reduces support resolution times by understanding customer intent from first contact. Specialized Claude agents handle different service areas, including account inquiries, disputes and loan applications, with Amazon Connect routing customers to the right agent immediately. The system analyzes previous interactions, account status and real-time sentiment to provide comprehensive context before conversations begin.

Omnichannel support maintains consistent experiences across all touchpoints. Bedrock's unified API deploys a single Claude agent that preserves context across channels: Amazon Lex for voice, Amazon Pinpoint for mobile and AWS AppSync for real-time synchronization. Customers can start a mortgage application on mobile and seamlessly continue with a branch representative, with Claude maintaining full context and preparing documentation via Textract.

Escalation and human-in-the-loop design ensures smooth handoffs for complex issues. Amazon Connect monitors sentiment and complexity scores, triggering escalation workflows when needed. Claude prepares detailed summaries including history, attempted resolutions, and recommendations, while Lambda functions enable real-time expert guidance without interrupting the customer experience.

Portfolio optimization and advisory

Not every client gets access to your best advisors. AI agents help you deliver institutional-quality advice at scale.

Personalized investment strategies combine market analysis with individual client circumstances. Claude agents integrate client data from CRM systems, market feeds via AWS Data Exchange and macroeconomic indicators to generate customized proposals. The system explains complex concepts in accessible language, automatically generating client-ready reports with clear reasoning, backtested performance, and risk scenarios.

Risk assessment and rebalancing becomes continuous and intelligent with 24/7 portfolio monitoring. Amazon EventBridge triggers Claude's analysis based on market events or client changes, evaluating thousands of scenarios using Monte Carlo simulations. Claude explains recommendations considering tax implications, transaction costs, and client constraints (stored in RDS) while maintaining audit trails in Amazon S3.

Compliance and regulatory adherence embeds directly into advisory agents. Bedrock's private endpoints deploy Claude agents fine-tuned on specific compliance policies. Pre-trade compliance validates proposed trades against mandates, restrictions, and risk limits while automatically generating required documentation including MiFID II reports.

Code modernization and legacy system migration

Your COBOL systems work—they're just impossible to maintain. AI agents help you modernize gradually without risking critical operations.

COBOL to modern language translation goes beyond syntax conversion to preserve business logic while modernizing architecture. Claude agents translate code while understanding embedded functionality, managing version control, and generating test suites.

Mainframe to cloud transitions accelerate through comprehensive ecosystem analysis. Claude agents work with AWS Mainframe Modernization service to map datasets to Amazon S3, convert legacy transaction systems into containerized microservices and transform batch jobs to workflows.

Documentation and knowledge preservation captures decades of embedded expertise. Claude analyzes existing documentation, code comments, and runbooks, synthesizing information into knowledge bases. Interactive sessions with experts generate technical documentation, architectural diagrams, and training materials, ensuring successful knowledge transfer across technology generations.

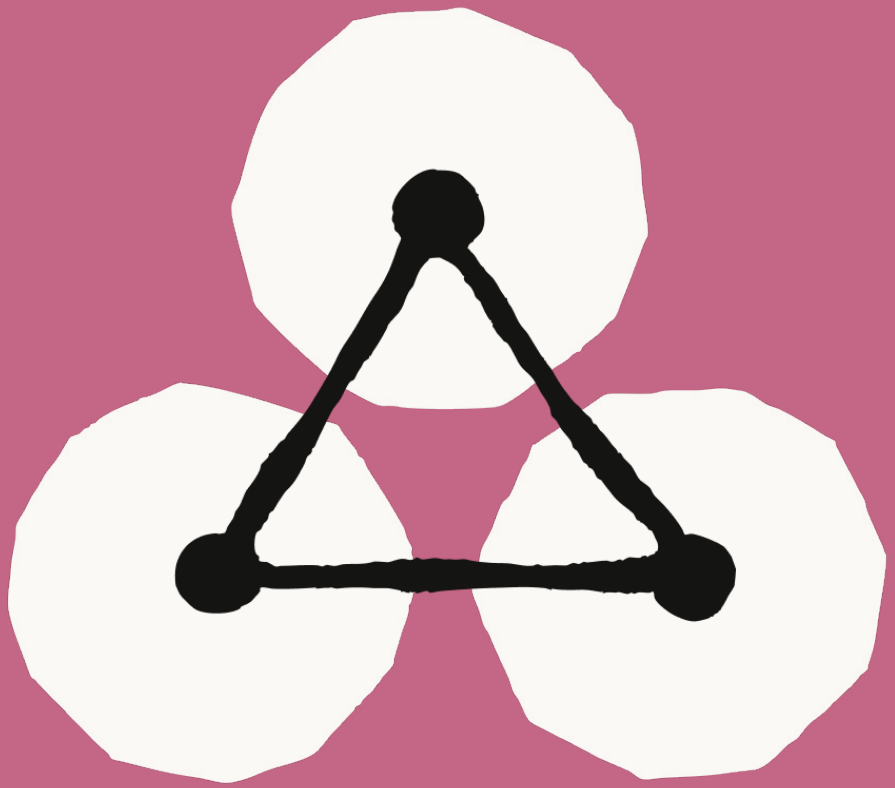

Chapter 4

Architecture patterns and agentic system design

Chapter 4

Architecture patterns and agentic system design

Financial institutions implementing Claude in Bedrock can choose from several proven architecture patterns, each designed to address different organizational needs and maturity levels. These patterns provide blueprints for deployment, scaling from simple proof-of-concept implementations to enterprise-wide AI ecosystems.

Single-agent deployment model

Perfect for your first production agent, this pattern focuses on one specific use case.

When to use this: Investment research support, compliance checking, or customer inquiry handling. Start here to prove value before scaling.

Architecture components

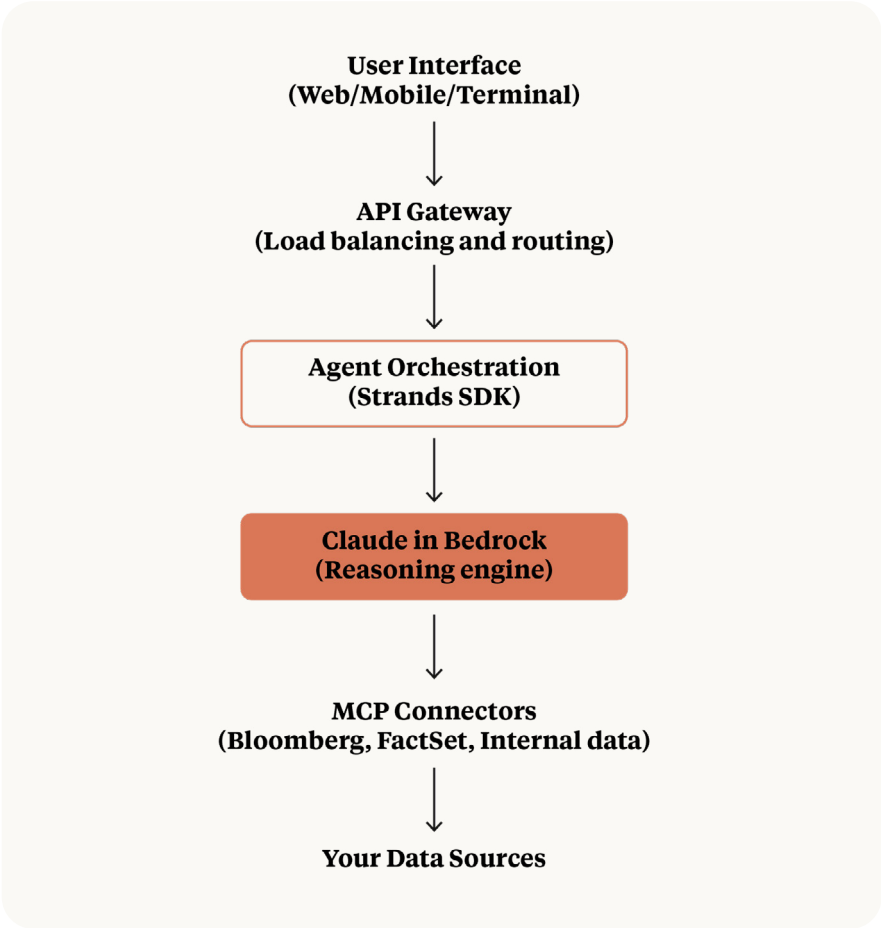

Scale up: multi-agent orchestration

The multi-agent orchestration pattern enables sophisticated workflows where specialized agents collaborate to solve complex financial problems, such as comprehensive due diligence or multi-asset portfolio optimization.

Architecture components

These reference architectures provide blueprints for financial institutions to deploy Claude-powered AI agents that are secure, scalable, and seamlessly integrated with their existing technology landscape, whether starting with focused single-agent deployments or scaling to sophisticated multi-agent ecosystems.

Access full implementation details for Claude in Bedrock, here.

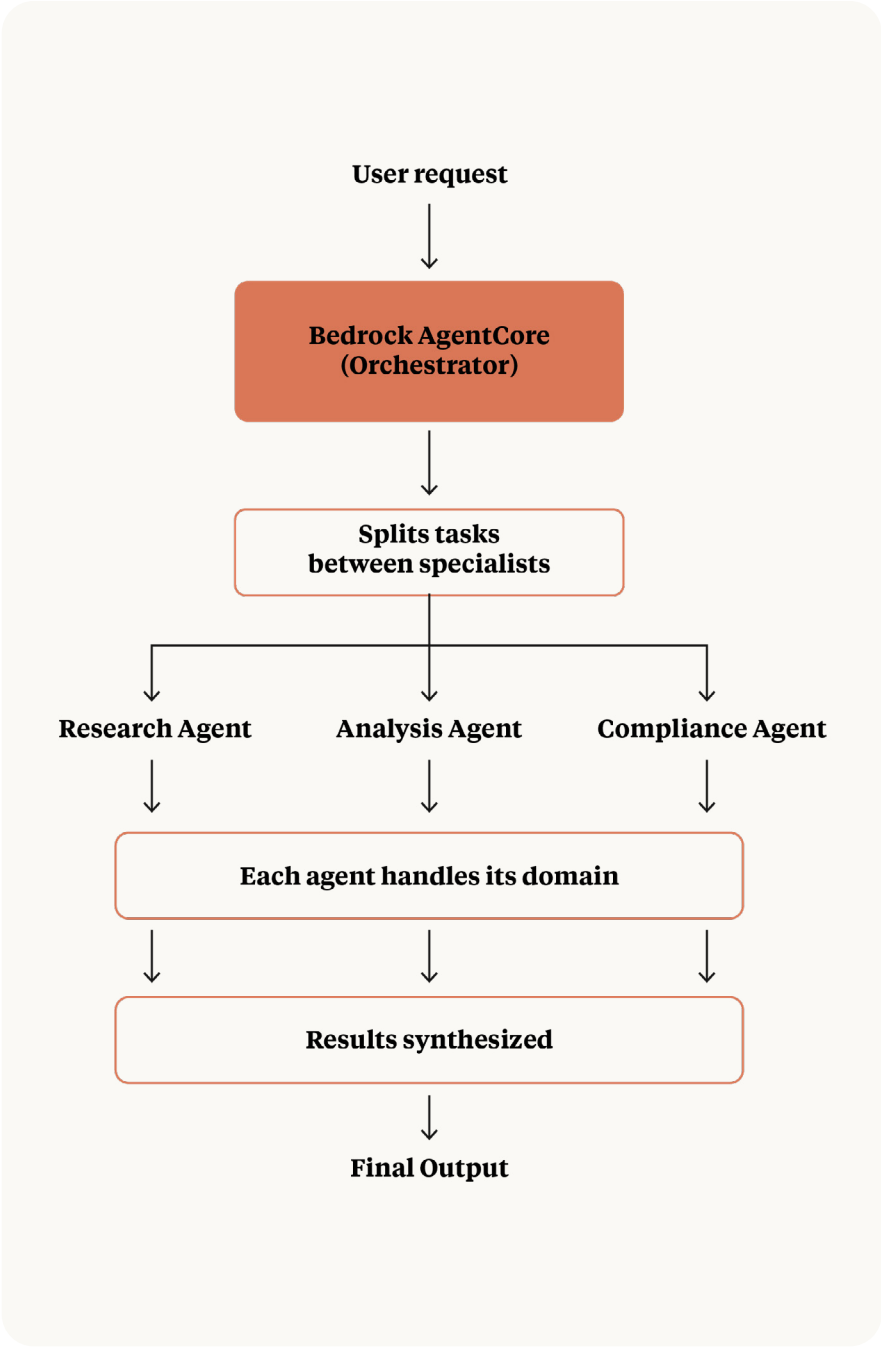

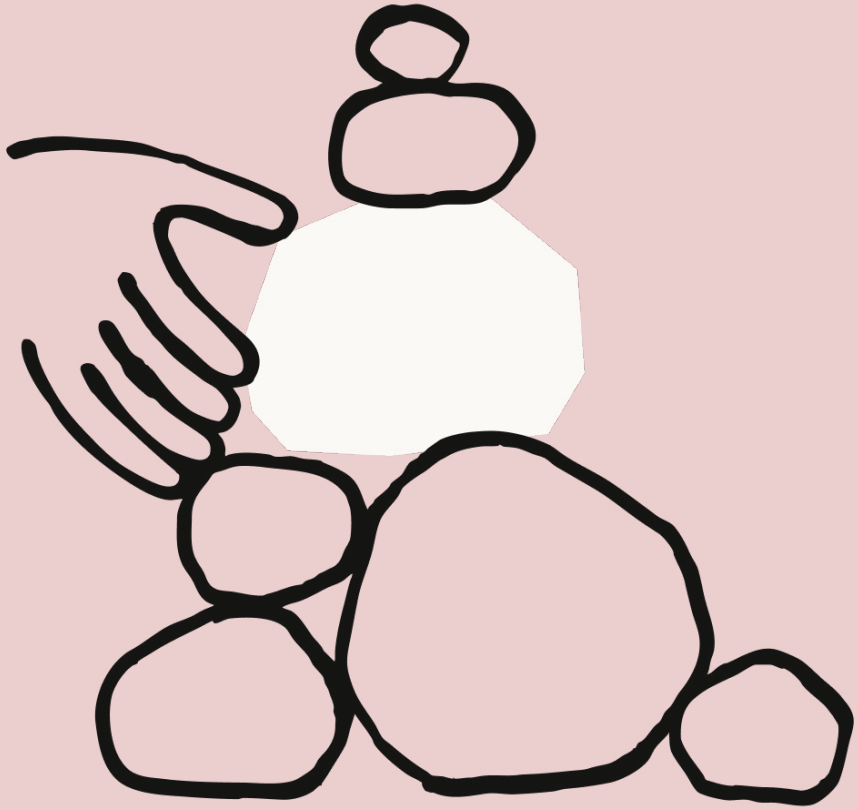

Chapter 5

Building your implementation roadmap, timelines, and team

Chapter 5

Building your implementation roadmap, timelines, and team

Implementing autonomous AI agents in financial services requires careful planning that addresses regulatory requirements, risk management, and operational resilience. This chapter outlines a structured approach for deploying Claude-powered agents on AWS within financial institutions.

Implementation roadmap

Phase 1: Prove value with a pilot

- **Pick your battle carefully:** Choose a painful but contained problem. Equity research support or tier-1 customer inquiries work well. Avoid anything touching core trading systems initially.
- **Build your coalition:** Partner with 2-3 senior users who'll champion the project. Their feedback and credibility matter more than perfect technology.
- **Define clear success metrics:** For instance, 85% accuracy on defined tasks, 50% time savings, positive user feedback. Make success measurable and achievable.

Phase 2: Expand strategically

- **Connect to more systems:** Add market data feeds, CRM systems, and risk databases. Each integration multiplies your agent's value.
- **Coordinate agents:** Progress from single agents to teams. Let research agents feed portfolio construction agents. Start simple workflows before attempting complex orchestration.
- **Build your AI team:** Combine technologists with business experts. The best agents come from teams that deeply understand both domains.

Phase 3: Scale across the enterprise

- **Deploy agent networks:** Enable complex workflows like automated compliance checks and continuous portfolio rebalancing.
- **Implement proper governance:** Align with model risk guidelines like SR 11-7. Document everything. Make decisions auditable.
- **Measure real impact:** Track operational errors, response times, and compliance scores. Show concrete ROI to secure continued investment.

Realistic timelines for your planning

In financial services, realistic implementation timelines are critical for success. Unlike other industries where teams can deploy and iterate quickly, financial institutions must navigate regulatory approvals, risk assessments and stakeholder alignment at each stage. Setting appropriate expectations upfront helps secure executive buy-in, allocate sufficient resources, and avoid the perception of failure when deployments take longer than anticipated.

To scope effectively, start by categorizing your use case along these dimensions:

- **Low-risk internal agent:** 4-5 months (research assistant, back-office automation)
- **Medium-risk customer-facing:** 6-9 months (service chatbot, account opening)
- **High-risk trading/advisory:** 9-12 months (portfolio management, trade execution)

See below for a timeline template you can leverage:

Implementation Timeline Template

Phase	Key Milestones	Expected Outcomes
Technical Foundation	• Bedrock setup in production VPC security controls per FFIEC guidelines • Claude model access and testing • MCP server deployment for 1-2 data sources	Secure infrastructure meeting compliance requirements
Pilot Development	• Use case documentation and approval • Agent development with domain experts • Testing with business users • Model validation documentation	Validated agent with documented performance metrics
Production Readiness	• Information security review • Model risk assessment (SR 11-7) • Business continuity planning • Controlled production release	Approved agent in production with monitoring
Scaling and Optimization	• Additional use case deployment • Cross-functional agent coordination • Performance tuning • Compliance reporting integration	Multiple agents supporting business processes

Building your team

The implementation timelines outlined above demand specialized teams that can operate effectively within financial services constraints. Success isn't just about having AI expertise; it's about assembling professionals who understand how to translate that expertise into compliant, production-ready systems that meet institutional standards.

Resourcing Suggestions

Role	FTE Allocation	Key Responsibilities	Required Experience
Core Technical Team			
AI/ML Engineer	1-2	Agent development, prompt optimization, testing	LLM experience, Python, some financial markets knowledge
Cloud Engineer	1	AWS infrastructure, security implementation	AWS certified, financial services cloud experience
Integration Developer	1-2	System connectors, API development	Capital markets systems, FIX protocol, REST APIs
Platform Engineer	1	Deployment automation, monitoring setup	DevSecOps, financial services compliance
Business Stakeholders			
Business Analyst	1-2	Requirements gathering	Front/middle office experience, process documentation
Product Owner	1	Stakeholder management, prioritization, user testing coordination	Trading/banking operations, agile delivery
Training Lead	0.5-1	User adoption, documentation	Financial services training, change management
Risk and Compliance			
Model Risk Manager	0.25	Model validation, ongoing monitoring	SR 11-7 expertise, quantitative background
Compliance Analyst	0.25	Regulatory review, policy updates	FINRA/SEC regulations, AI governance
Information Security	0.25	Security assessments, access reviews	Financial services security, cloud security

Chapter 6

Financial services agents in the real world

Chapter 6

Financial services agents in the real world

Transforming the world's largest sovereign wealth fund with AI

Challenge and solution overview

When Norwegian Bank Investment Management (NBIM), manager of Norway's $1.8 trillion sovereign wealth fund, faced the challenge of scaling operations without expanding their 700-person workforce, they turned to AI for a solution. With the fund trading around the clock across global markets and funding 20% of Norway's state budget, NBIM needed to dramatically increase efficiency while maintaining their high standards for investment analysis and decision-making.

NBIM partnered with Anthropic and AWS to deploy Claude across their entire organization, creating a comprehensive AI transformation program. This enterprise-wide implementation uses Claude for Enterprise alongside custom API integrations to enhance everything from investment analysis to administrative tasks, with the ambitious goal of achieving 20% efficiency gains across all operations.

Implementation journey

NBIM built a multi-layered AI strategy centered on organizational change management and technical integration. Their implementation included Claude for Enterprise deployed to all 700 employees across laptops, phones, and tablets; custom MCP (Model Context Protocol) integrations connecting Claude to internal data sources including their Snowflake data warehouse; mandatory 30-minute training modules developed with Anthropic for all staff; an AI ambassador network of 20 volunteers driving adoption across departments; and complementary tools including Cursor for developers and Microsoft Copilot for Office tasks.

The organization took a "learning by doing" approach, encouraging employees to create projects and share best practices. They established clear principles around data confidentiality levels and security assessments while ensuring AI tools could be used safely with sensitive financial data. CEO engagement proved critical, with leadership regularly asking employees "How have you used AI today?" to drive cultural adoption.

Results and ROI

NBIM's AI transformation delivered measurable results within the first year:

- 20% efficiency improvement achieved by June
- 350,000+ conversations in Claude despite having only 700 users
- 65,000+ projects and 172,000+ files created across the organization, from investment analysis to HR support
- Transformed workflows enabling analysts to analyze more companies, increase analysis frequency, and improve quality
- Organization-wide adoption with employees using Claude for everything from complex financial analysis to coding and personal productivity

The impact extends beyond metrics, with employees reporting they can focus on more strategic, enjoyable work rather than tedious tasks. As their Head of Service Provider Management noted, “I can now move on from tedious requirements writing to improving our processes, following up more on vendors, or doing greater budgeting.” The cultural transformation has been so complete that employees now bring Claude home, using it for personal tasks and even helping their children with schoolwork.

Learn more here.

Getting started

Throughout this guide, we've shown how teams like yours are using AI agents to solve real problems—automating painful workflows, delighting customers, and staying ahead of compliance requirements. The combination of Claude's intelligence and AWS's infrastructure gives you everything needed to join them.

The future belongs to teams that start building today.

Ready to move from endless planning to agents that ship?

Contact Anthropic at sales@anthropic.com or AWS Financial Services.

Learn more: https://claude.com/solutions/financial-services

Resources to get started:

Anthropic Documentation: Complete guides at docs.anthropic.com

Anthropic Academy: Level up your team at academy.anthropic.com

Building effective agents: Engineering best practices at anthropic.com/research/building-effective-agents

AWS Financial Services: Latest innovations at aws.amazon.com/blogs/industries/category/industries/financial-services/

Amazon Bedrock: Deployment guides at docs.aws.amazon.com/bedrock/

ANTHROP\C | aws

https://claude.com/partners/amazon-bedrock