



AEGIS RESEARCH SERIES · CASE STUDY

Claude Inside

Architecture, connectors, and governance in a regulated financial institution – a technical deep dive on the Model Context Protocol (MCP) connector layer and how it maps onto JPMorgan’s enterprise architecture.

ANTHROPIC CLAUDE PARTNER – ARIANA DIGITAL LLC

PREPARED
July 2026

SCOPE
Public-record analysis

PRACTICE
AEGIS Governance Advisory

PUBLISHER
Ariana Digital LLC

CONTENTS

Table of Contents

Abstract & Methodology**1 Organizational Context****2 WHAT – The Deployment Landscape****3 WHY – Strategic Rationale & Measured Outcomes****4 HOW – Implementation Model & Governance****5 Deep Dive – Connectors & Enterprise Architecture Fit**

5.1 The integration problem MCP was built to solve

5.2 MCP architecture fundamentals

5.3 Mapping the connector layer onto JPMorgan's stack

5.4 Deployment shapes

5.5 The broader 2026 reference architecture

5.6 Security posture of the connector layer

5.7 Beyond JPMorgan – Claude and enterprise platforms

6 Risk, Governance & Open Questions**7 Roadmap – What's Accelerating vs. Holding Back****8 Key Metrics Summary****9 Limitations of This Case Study****References****Appendix A – The Ten Claude for Financial Services Agent Templates**

ABSTRACT

Claude Inside JPMorgan Chase

JPMorgan Chase operates the largest publicly documented enterprise deployment of Claude in financial services. Rather than exposing Claude directly to its workforce, the bank built a proprietary abstraction layer — **LLM Suite** — that routes ~250,000 employees across a model-agnostic interface spanning Anthropic and OpenAI models, sitting on top of two internal platforms (**JADE** for data, **OmniAI** for MLOps) that predate the generative-AI era by nearly a decade ^{[1][3][4]}. Layered above that general-purpose tool is a narrower, higher-stakes deployment: **Claude for Financial Services**, a set of purpose-built agents running on Claude Opus 4.7 and wired into market-data and enterprise-data providers via the **Model Context Protocol (MCP)** ^{[12][14]}.

This case study documents what has been deployed, why the bank made the investment, how the rollout was governed, and — as the primary technical contribution — how the MCP connector layer is architected and where it sits relative to JPMorgan's existing data and ML infrastructure.

METHODOLOGY & SOURCE NOTE

This is a research case study built entirely from public disclosures — earnings calls, executive interviews, Anthropic product announcements, and financial/trade press — not from JPMorgan or Anthropic internal documentation. Adoption metrics (efficiency gains, cost savings, headcount) are management-reported figures cited by journalists, not independently audited numbers. Where a claim is single-sourced or comes from a company's own promotional material, that is flagged inline. Readers using this document for investment, procurement, or competitive-strategy decisions should treat the quantitative claims as directionally credible, not verified.

1

ORGANIZATIONAL CONTEXT

Organizational Context

JPMorgan's AI program did not start with generative AI. Dimon has stated the bank began using AI in some form as far back as 2012, well predating the ChatGPT-driven wave ^[12]. That matters architecturally: LLM Suite and Claude for Financial Services were not built on greenfield infrastructure. They were bolted onto, and had to interoperate with, a decade-plus-old data and machine-learning stack — the central design constraint that shapes everything in Section 5.

By 2026, JPMorgan has topped the Evident AI maturity index for four consecutive years and carries an annual technology budget near \$18–19.8 billion, with AI-specific spend estimated at up to \$2 billion a year ^{[1][10]}. The bank's own executives frame the destination state as a full "agentic enterprise": Chief Analytics Officer Derek Waldron has described a target where every employee has a personalized AI assistant, every internal process is agent-powered, and every client interaction includes an AI concierge ^{[5][8]}.

Every employee has a personalized AI assistant, every internal process is agent-powered, every client interaction includes an AI concierge.

DEREK WALDRON, CHIEF ANALYTICS OFFICER, JPMORGAN CHASE — PARAPHRASED FROM ^{[5][8]}

2 WHAT

The Deployment Landscape

Scale. LLM Suite onboarded over 200,000 employees within eight months of its July 2024 launch, growing past 230,000–250,000 by mid-2026, with roughly half of licensed users active daily ^{[2][5][7]}. Claude for Financial Services, by contrast, launched in July 2025 and expanded in May 2026 with a narrower, higher-value footprint – reported in production at JPMorganChase, Goldman Sachs, Citi, AIG, and Visa ^{[11][12]}.

TABLE 1 – WHERE CLAUDE SHOWS UP ACROSS THE BANK

LAYER	SYSTEM	FUNCTION	MODEL(S)
General employee tooling	LLM Suite	Firm-wide copilot: drafting, summarization, decks, performance reviews	Anthropic Claude + OpenAI, selectable per user ^{[4][5]}
Financial analysis agents	Claude for Financial Services	10 pre-built agents incl. Earnings Reviewer, Market Researcher, Model Builder	Claude Opus 4.7 ^{[12][14]}
Software engineering	Claude Code	Dashboard / data-analysis generation, code authoring	Claude (used personally by CEO Jamie Dimon) ^{[11][12]}
Code review	PRBuddy	Internal PR / code-review automation built on LLM Suite	Claude via LLM Suite ^[6]
Wealth management	Connect Coach	Real-time advisor insight generation	Claude via LLM Suite ^{[4][10]}
Legal	COiN	Contract / document analysis at scale	Not publicly confirmed as Claude-specific; part of the broader AI stack ^[1]
Call center	EVEE Intelligent Q&A	Real-time agent-assist for customer-service reps	Model not publicly disclosed ^[2]
Cybersecurity (restricted)	Project Glasswing / Claude Mythos	Vulnerability discovery & patching ahead of adversarial exploitation, invitation-only access	Claude Mythos ^[16]

EXHIBIT 1 – LAYERED IMPLEMENTATION STACK



3 WHY

Strategic Rationale & Measured Outcomes

TABLE 2 – STRATEGIC RATIONALE & MEASURED OUTCOMES

METRIC	REPORTED VALUE	SOURCE
General LLM Suite efficiency gain	30–40%	Forbes, JPMorgan-reported ^[1]
Engineering / coding assistant efficiency gain	10–20%	Tearsheet, Emerj ^{[2][4]}
Annual AI-driven cost savings (through 2025)	~\$2 billion / year	Forbes, citing CEO Jamie Dimon ^[1]
Annual value from AI in wealth / asset management	Up to \$1.5 billion	Emerj ^[4]
Legal work automated via COiN	360,000 hours	Forbes ^[1]
Time saved per employee per week (LLM Suite)	3–6 hours	The Digital Banker ^[7]
Insurance underwriting accuracy improvement (AIG, comparable deployment)	75% → 90%, 5x faster review	Anthropic-adjacent press ^[15]
Deck generation time (investment banking)	Hours → ~30 seconds	AI Magazine, European Financial Review ^{[5][8]}

Four Strategic Drivers

O1

Competitive positioning

Anthropic has explicitly marketed Claude as becoming “the operating layer for Wall Street,” using JPMorgan as its flagship reference client alongside Goldman Sachs, Citi, AIG, and Visa ^[11].

O2

Cost & margin pressure

With a near-\$20 billion annual tech budget, even single-digit productivity gains across 250,000 employees compound into material savings — the reported \$2 billion/year figure reflects this arithmetic ^{[1][10]}.

O3

Workforce restructuring

Consumer-banking leadership has told investors operations headcount could fall at least 10% over five years as a direct consequence of AI deployment — a structural labor-model shift, not a productivity tool alone ^[5].

O4

Cybersecurity value

Access to Claude Mythos was granted specifically to find and patch vulnerabilities before adversaries could exploit them — linked by journalists to explicit executive concern about AI-augmented cyberattack risk ^[16].

4 HOW

Implementation Model & Governance

INTERNAL-FIRST SEQUENCING

JPMorgan's rollout followed a deliberate three-pillar strategy: (a) heavy, multi-year investment in foundational data/ML infrastructure before any generative-AI layer; (b) an "internal-first" philosophy that proves value and builds AI literacy on employee-facing tools before any client-facing exposure; and (c) proprietary applied research to build durable differentiation rather than renting a vendor's off-the-shelf product ^[3]. This sequencing explains why the bank built LLM Suite as a wrapper instead of simply issuing ChatGPT/Claude seats — a wrapper gives the bank a governance chokepoint.

MODEL-AGNOSTIC ROUTING AS RISK MANAGEMENT

Employees pick a model from a dropdown inside LLM Suite; Claude and OpenAI models sit side by side ^{[4][5]}. This buys JPMorgan negotiating leverage and vendor-outage resilience — and, as Section 6 shows, turned out to be operationally important when Anthropic-specific licensing restrictions hit a specific geography.

RISK-TIERED AUTONOMY

JPMorgan's CIO organization does not apply uniform autonomy to AI agents. In HR contexts, a human retains broader visibility into employee data than an agent is granted; in software engineering, agents are given more latitude because a validation layer exists to catch and correct errors before code ships ^[10].

Autonomy scales with the cost of an undetected error, not with technical capability.

GOVERNANCE PATTERN DOCUMENTED IN THIS STUDY – WORTH GENERALIZING BEYOND JPMORGAN

CONTRACTUAL DATA PROTECTION AS A FLOOR REQUIREMENT

Anthropic's enterprise commitment — that customer data is not used to train Anthropic's models by default — is treated as a baseline, non-negotiable term for a bank under OCC/Federal Reserve supervision, not a differentiator ^[14].

5

DEEP DIVE

Connectors & Enterprise Architecture Fit

This section is the technical core of the case study: how Claude’s tool-use layer — implemented via the **Model Context Protocol (MCP)** — is architected, and how it maps onto JPMorgan’s pre-existing enterprise data and ML platforms.

5.1 The Integration Problem MCP Was Built to Solve

Before MCP, connecting N AI applications to M enterprise data sources required custom point-to-point integration for every pairing — the so-called **$N \times M$ problem**. Ten AI tools needing access to a hundred internal systems implies up to a thousand bespoke integrations, each with its own auth model, error handling, and maintenance burden [22][23]. Anthropic introduced MCP as an open standard in November 2024 specifically to collapse $N \times M$ into $N + M$: build one MCP client implementation per AI application and one MCP server per data source, and any compliant client can talk to any compliant server [20][22]. In December 2025, Anthropic donated MCP’s stewardship to the **Agentic AI Foundation (AAIF)**, a directed fund under the Linux Foundation co-founded with Block and OpenAI — a governance move intended to keep the protocol vendor-neutral as it becomes foundational infrastructure, comparable to how Kubernetes or Node.js is governed [22][35].

5.2 MCP Architecture Fundamentals

MCP defines three architectural roles [22][30][36]:

- **Host** — the AI application itself (Claude for Financial Services, Claude Code, Claude Desktop, LLM Suite’s underlying orchestration layer).
- **Client** — a 1:1 connector instance inside the host that manages a single stateful session with one server.
- **Server** — a lightweight process exposing a specific system’s capabilities in a standard, self-describing format.

TABLE 3 – THE THREE MCP PRIMITIVES

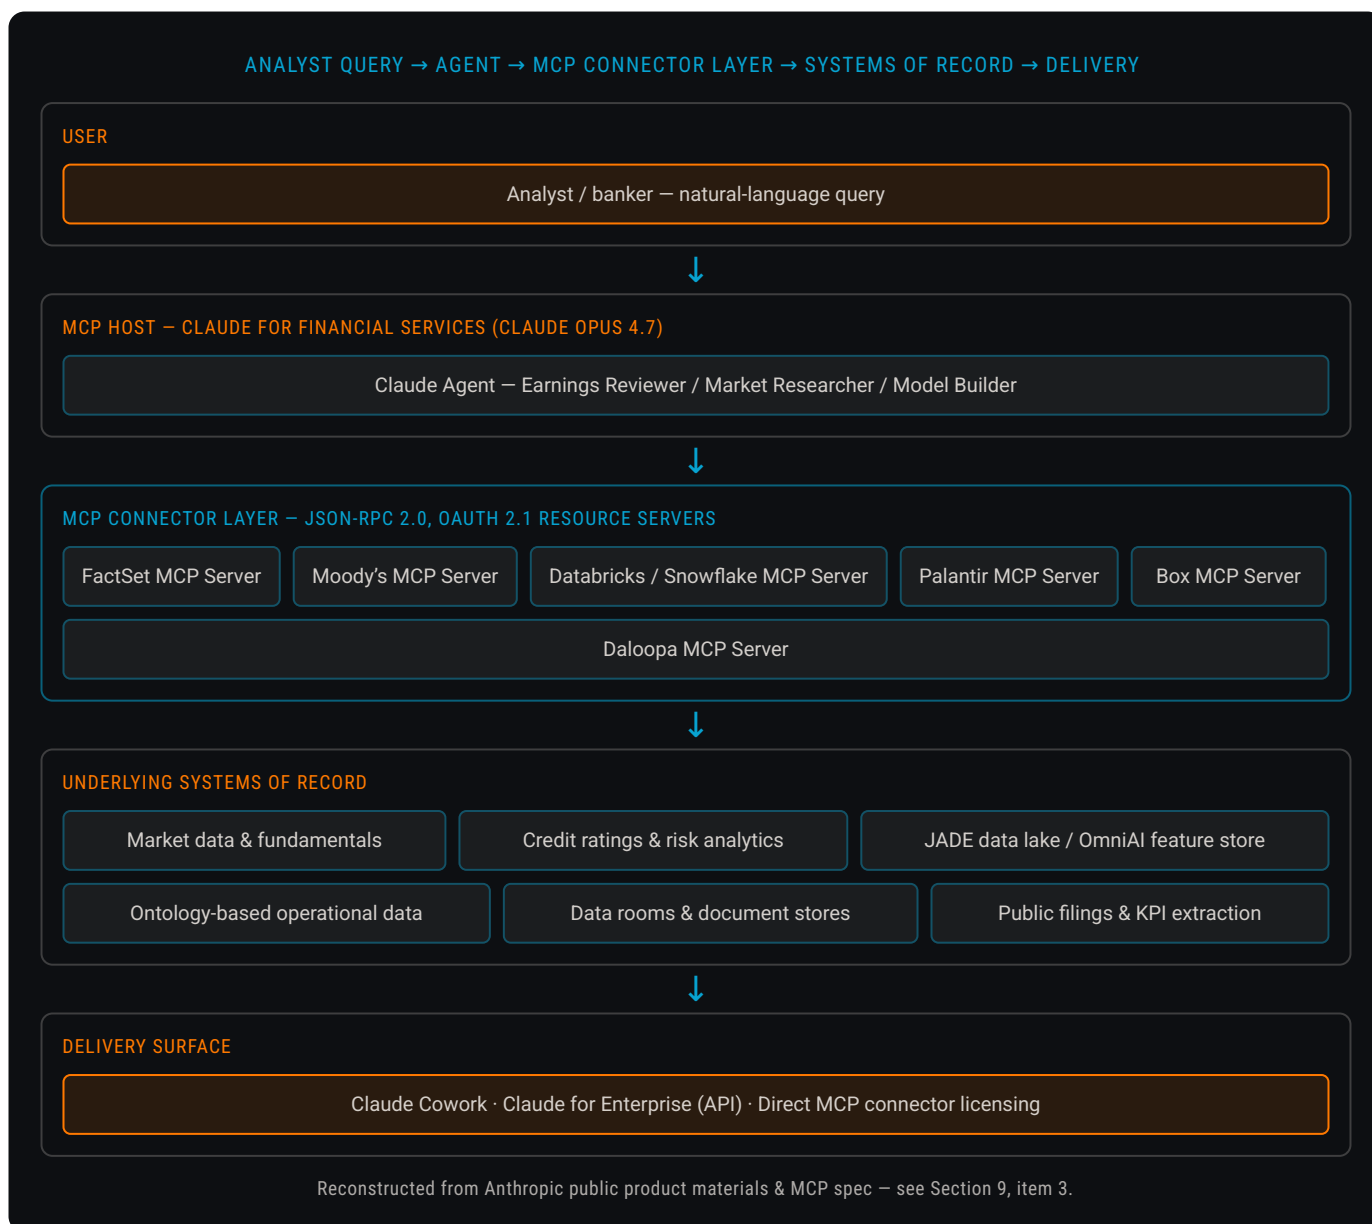
PRIMITIVE	WHAT IT IS	JPMORGAN-RELEVANT EXAMPLE
Tools	Functions the model can invoke	Pull FactSet consensus estimates for a named ticker
Resources	Data the model can read	A specific data-room document set in Box for a live deal
Prompts	Predefined workflow templates	A standardized DCF or LBO build sequence

Communication runs over **JSON-RPC 2.0**, the same message family that underlies the Language Server Protocol used by code editors ^{[22][32]}. Two transport modes matter for an enterprise deployment: **stdio** (local, in-process – Claude Code and Claude Desktop connecting to local tools) and **Streamable HTTP** (the November 2025 spec update, replacing the older SSE transport, allowing MCP servers to run as scalable remote services behind a load balancer) ^[30]. Enterprise-grade MCP servers authenticate as **OAuth 2.1 Resource Servers**, advertising their authorization endpoint via a .well-known discovery path rather than relying on static API keys ^[30].

5.3 Mapping the Connector Layer onto JPMorgan's Stack

Anthropic's own materials confirm Claude for Financial Services ships with pre-built MCP connectors into FactSet (equity fundamentals and consensus estimates), Moody's (credit ratings and risk analytics), Databricks and Snowflake (internal analytics warehouses), Palantir (ontology-based enterprise data integration), Box (secure document/data-room access), and Daloopa (structured KPI extraction from public filings) – alongside a full Microsoft 365 integration spanning Excel, PowerPoint, Word, and Outlook ^{[14][15]}.

EXHIBIT 2 – CONNECTOR ARCHITECTURE & DATA FLOW



THE CRITICAL ARCHITECTURAL POINT

The MCP connector layer is **additive, not substitutive**. It does not replace JADE (JPMorgan’s internal data-governance and pipeline layer) or OmniAI (its ML development/deployment platform) — it sits alongside them as another consumer of the same underlying systems of record, ideally routed through the same access-control policies those platforms already enforce. Practitioners increasingly call this a “**governed MCP server**” pattern: rather than building parallel, AI-specific governance infrastructure, the MCP server extends the bank’s existing data-governance policy to AI interactions ^[24]. For an institution under continuous regulatory examination, reusing an existing, audited control plane is arguably more valuable than any single model-capability upgrade — a new AI use case doesn’t require a net-new security review from scratch.

5.4 Deployment Shapes

Public reporting describes three distinct packaging models for how the Claude for Financial Services agents (and their MCP connectors) reach end users ^[15]:

- **Claude Cework** — a collaborative workspace where the ten agent templates are distributed as installable plugins through an internal marketplace; an IT administrator enables licensed templates, and analysts invoke them conversationally, iterating and reviewing before export.
- **Claude for Enterprise (API-based)** — direct programmatic access with expanded usage limits, for teams building custom workflows (e.g., modernizing trading systems, Monte Carlo risk simulations) rather than using pre-packaged templates. Distributed via AWS Marketplace with consolidated billing.
- **Direct MCP connector licensing** — data-provider connectors (FactSet, Moody's, etc.) require separate commercial agreements contracted directly with each data vendor; MCP standardizes how Claude talks to these systems, not the underlying data-licensing relationship or cost.

COST-MODELING NOTE

MCP reduces integration engineering cost, not data-licensing cost. A bank still pays FactSet for FactSet data; MCP just means it doesn't have to build and maintain a bespoke FactSet-to-LLM bridge to consume it inside a Claude workflow.

5.5 Where This Sits in the Broader 2026 Reference Architecture

Industry analysis published in the first half of 2026 converges on a common enterprise pattern: organizations are standardizing on a single “core orchestrator” (LLM Suite for general use; Claude for Financial Services as a specialized orchestrator for high-value workflows), with MCP as the standard mechanism by which that orchestrator reaches every other AI-relevant system ^[33]. A commonly cited benchmark places roughly two-thirds of enterprises still in “pilot phase” for this architecture, with only a small minority “fully scaled” ^[33] — which would place JPMorgan's rollout toward the leading edge of adoption maturity, not the median.

The MCP roadmap itself is still maturing in ways directly relevant to a deployment at JPMorgan's scale: current MCP servers must maintain session state, which constrains horizontal scaling behind load balancers; the protocol's 2026 roadmap standardizes stateless operation, session resumption, and migration so server restarts and scale-out events become transparent to connected clients ^[30]. For market-critical, latency-sensitive workflows, this statelessness work determines whether the connector layer can meet the same reliability SLAs the bank already demands of core trading and payments infrastructure.

5.6 Security Posture of the Connector Layer

Because an MCP server sits directly between a language model and a live enterprise system, its security model is the connector architecture's single highest-leverage risk surface. Best-practice guidance converging across 2026

enterprise-security literature emphasizes: MCP is protocol-neutral on governance – it standardizes how an AI tool and a data system talk, not what is permitted once they’re talking – so the security burden sits entirely with how the bank configures its MCP servers, not with the protocol itself [24]. For JPMorgan, this shows up as the layered governance already noted in Section 4: default no-training-on-client-data terms, OAuth 2.1 token lifecycle management (critical because agentic workflows run asynchronously), and geography/entity-level gating enforced above the connector layer entirely (Section 6).

5.7 Beyond JPMorgan — How Claude Integrates with Enterprise Platforms

Generally

Sections 5.3–5.6 covered JPMorgan’s named connector stack. That is one instance of a broader, product-wide pattern. Any enterprise adopting Claude integrates with its internal platforms through two distinct layers that are easy to conflate but function differently: an **application connector layer** (what Claude can reach) and an **identity/administrative layer** (what Claude is allowed to do once it gets there).

TABLE 4 – APPLICATION CONNECTOR LAYER

PLATFORM	INTEGRATION MODEL
Microsoft 365	First-party MCP connector; requires one-time Microsoft Entra Global Administrator consent per tenant
Slack	Bidirectional first-party MCP integration; “Claude Tag” lets employees @-mention Claude directly in a channel with scoped permissions
Salesforce	Deepest integration of the group — native embedding, not a bolt-on connector. Claude is a preferred model inside Agentforce; Anthropic is described as “the first LLM provider fully integrated within the Salesforce trust boundary” — inference stays inside Salesforce’s VPC, routed through the Einstein Trust Layer with zero data retention on prompts [34]
Atlassian (Jira, Confluence, Compass)	Atlassian’s own remote MCP server, with Claude as its first named official partner; GA February 2026
ServiceNow, Workday, SAP, NetSuite, similar ITSM/ERP	No first-party Anthropic connector found in current sourcing; reached via third-party MCP bridge vendors (e.g., CData)
Google Workspace	Inconsistent coverage as of early-to-mid 2026: Claude Cowork did not yet support it as of February 2026, while it is listed elsewhere as a supported SSO identity provider

A DISCIPLINE WORTH ADOPTING

The total “connector count” Anthropic or third parties cite varies by an order of magnitude across sources (roughly 38 curated connectors inside Cowork’s UI in one account versus 400+ in the broader public MCP directory in another) [36]. That gap is almost certainly a counting-methodology difference – curated first-party list versus the full open ecosystem – not a factual contradiction. “How many systems does Claude connect to” is not a number to cite without specifying which registry is being counted.

TABLE 5 – IDENTITY & ADMINISTRATIVE LAYER

CONTROL	WHAT IT DOES
SSO	SAML 2.0 and OIDC, provisioned through WorkOS; supports Okta, Microsoft Entra ID, Google Workspace [29][37]
SCIM / JIT provisioning	Automated user creation, role assignment, and deprovisioning synced to the corporate directory; removing someone from the IdP revokes Claude access across web, desktop, and CLI [37][38]
RBAC	Three-tier default model (Primary Owner / Admin / Member); per-connector, group-level RBAC scoping is in beta as of early 2026 [38]
Audit logs	Metadata only (who, action, timestamp, device) – not conversation content. Retained 180 days by default for Enterprise, exportable to SIEM. Reading prompt/response content requires the separate Compliance API [38]
Data residency	Workspace-level workspace_geo setting, immutable once created; “us” and “global” region options [37]
Zero Data Retention (ZDR)	Optional contractual addendum for Enterprise API traffic; not automatically extended to web UI or beta features unless separately contracted [37]

WHY THIS MATTERS FOR THE JPMORGAN CASE SPECIFICALLY

JPMorgan did not adopt Anthropic’s own admin console, connector directory, and RBAC model wholesale – it built LLM Suite as a proprietary substitute that reimplements the same function (identity gating, audit, data governance) on the bank’s own infrastructure, and only lets Claude’s native connector/MCP layer operate underneath that wrapper for the financial-services agent workflows specifically. That is a deliberate, expensive choice, not a default. For an enterprise without JPMorgan’s balance sheet or regulatory profile, adopting Anthropic’s first-party admin layer directly – SSO/SCIM plus the connector directory – is the far more common and far cheaper path to the same governance outcome. The tradeoff is control: JPMorgan’s approach enforces bank-specific policy (e.g., the Hong Kong geographic gate in Section 6) at a layer it fully owns.

6 RISK & GOVERNANCE

Risk, Governance & Open Questions

THE HONG KONG ACCESS REMOVAL (JUNE 2026)

In June 2026, JPMorgan removed Claude from the internal dropdown list of approved models available to staff in Hong Kong, following an earlier, similar move by Goldman Sachs in April 2026 ^{[17][18][19]}. Reporting attributes this to language in Anthropic’s usage-licensing terms restricting use across Greater China, rather than to any performance or security concern with the model itself ^[17].

Model access is not purely an IT provisioning decision — it is a jurisdiction-by-jurisdiction legal determination that can change independent of the technical deployment.

GOVERNANCE IMPLICATION DRAWN FROM THE HONG KONG ACCESS-REMOVAL EVENT

An architecture that hard-codes a single model choice into workflow logic is more fragile to this kind of change than one that preserves model-agnostic routing at the orchestrator layer — precisely the design decision LLM Suite made from the outset ^{[4][5]}.

EXPORT-CONTROL EXPOSURE

Separately, U.S. export-control action in mid-2026 required Anthropic to suspend international availability of its most advanced Mythos-class models, a restriction distinct from, but concurrent with, the Hong Kong licensing story ^[17]. JPMorgan’s Project Glasswing access to Claude Mythos for cybersecurity work is explicitly described as limited-access, invitation-only distribution to a small set of trusted organizations — a deliberate containment strategy given the model’s demonstrated ability to find previously unknown software vulnerabilities ^[16].

OPEN QUESTIONS NOT RESOLVED BY PUBLIC SOURCING

- Whether JPMorgan’s MCP servers are self-hosted internally (behind the bank’s own network boundary) or consumed as managed/remote MCP endpoints — this materially changes the audit and liability posture and is not disclosed in any source reviewed.
- The actual per-seat or per-connector commercial terms of the JPMorgan–Anthropic licensing agreement remain confidential; only the existence and general nature of licensing-term restrictions (the Hong Kong clause) have surfaced publicly.
- Whether COiN and EVEC run on Claude specifically or on a different model within LLM Suite’s multi-model roster — public sourcing does not confirm the underlying model for either system.

7 ROADMAP

What's Accelerating vs. What's Holding Back

Every accelerant below has a corresponding constraint pulling against it. This is framed deliberately as a tension rather than a scorecard, because that is how JPMorgan's own executives describe it in their own public statements — the sourcing below leans heavily on the bank's own research arm, its own CIO, and its own payments executives, not outside critique.

ACCELERATING

Infrastructure debt largely paid off. CFO Jeremy Barnum: modernization investment “has peaked”; cloud workloads moved from 50% to 65% in a single year [\[44\]](#).

OmniAI decouples “deploy AI” from “rebuild core systems.” A cloud-native integration layer lets JPMorgan ship AI features without a wholesale infrastructure rebuild [\[44\]](#).

Proven evaluation methodology exists internally. The Private Bank's “Ask David” system improved from ~50% to over 90% accuracy through evaluation-driven development with human-in-the-loop oversight [\[44\]](#).

Aggressive, quantified expansion target. Roughly 450 AI use cases in production, targeting ~1,000 by end of 2026 [\[4\]](#).

Vendor-side capability compounding independently. Claude for Financial Services, Claude Managed Agents, and native Microsoft 365 integration all shipped within the study window [\[12\]\[14\]\[39\]](#).

Cost economics aren't (yet) the brake. J.P. Morgan Payments' CDAO Zack Anderson: the firm isn't worried about escalating AI token costs; approach is strongly business-case-driven [\[41\]\[45\]](#).

HOLDING BACK

The bank's own research names the barriers. Security, regulatory compliance, and legacy-system integration remain immediate barriers — per J.P. Morgan Asset Management's own research [\[42\]](#).

Agentic treasury is explicitly early-stage. Data foundations, clear policies, and governance infrastructure are prerequisites not yet fully solved [\[41\]](#).

Delegation-threshold governance is unresolved. What the firm is willing to delegate to a machine, at what thresholds, under what conditions — full auditability is a precondition not yet fully built [\[41\]](#).

Security had to be rebuilt around agents, not models. A “lethal trifecta” framing: untrusted-input processing + sensitive-data access + authority to act externally requires safeguards that scale with compounding risk [\[43\]](#).

Organizational absorption is the bottleneck. CIO Lori Beer: the constraint isn't model capability anymore; it's getting hundreds of thousands of employees and legacy workflows to actually change [\[10\]](#).

Full automation remains the exception. 87% of organizations have implemented some treasury automation, but only 39% describe systems as fully automated [\[47\]](#).

NET READ

The accelerants and constraints above are mostly the same underlying fact, viewed from different vantage points. JPMorgan has the infrastructure, the capital, the executive mandate, and the vendor ecosystem to move faster than almost any peer institution. What it does not yet have — by its own executives' own public statements — is a fully solved answer to the governance question underneath every accelerant: at what threshold does a human stop being required in the loop, and how is that decision audited once made. That is a policy and organizational-change problem, not primarily a model-capability problem.

8

SCORECARD

Key Metrics Summary



TABLE 6 – FULL METRICS DETAIL

CATEGORY	METRIC	VALUE
Reach	Employees onboarded to LLM Suite	~230,000–250,000 [2][5][7]
Reach	Daily active usage rate	~50% of licensed users [5]
Productivity	General efficiency gain (LLM Suite)	30–40% [1]
Productivity	Engineering efficiency gain	10–20% [2][4]
Financial	Estimated annual AI-driven savings	~\$2B/year [1]
Financial	Annual tech / AI budget	~\$18–20B total tech; up to \$2B AI-specific [1][10]
Vendor breadth	Financial-services deployments citing Claude in production	JPMorganChase, Goldman Sachs, Citi, AIG, Visa [11]
Connector breadth	Named MCP-connected data providers	FactSet, Moody's, Databricks, Snowflake, Palantir, Box, Daloopa [14]
Governance event	Geography-restricted access	Hong Kong (June 2026) [17]

9

LIMITATIONS

Limitations of This Case Study

1. All quantitative outcomes are self-reported by JPMorgan or relayed secondhand through press interviews with bank executives – there is no independently audited ROI study cited in any source reviewed.
2. Several systems (COiN, EVEE) are described in public sourcing as part of JPMorgan’s broader AI stack without confirmation that Claude specifically, as opposed to another model in the LLM Suite roster, powers them.
3. Architecture details in Section 5 combine (a) Anthropic’s public product documentation for Claude for Financial Services, (b) general MCP protocol specification and industry commentary, and (c) inference about how these are likely combined at JPMorgan specifically. No JPMorgan technical architecture document was located publicly; Section 5.3’s mapping onto JADE/OmniAI is an informed synthesis, not a disclosed architecture diagram from the bank.
4. This document reflects the public record as of late June / early July 2026. The Hong Kong access story in particular is recent and may evolve as licensing terms are renegotiated or as U.S. export-control policy shifts.
5. Section 5.7’s product-wide integration survey is the most perishable part of this document. Anthropic ships connector and admin-console changes on a roughly monthly cadence; the Google Workspace coverage gap and exact Slack write-permission scope should be re-checked against current Anthropic documentation before being relied upon for a procurement decision.

INDEPENDENCE STATEMENT

This document is an independent research synthesis. It is not affiliated with, endorsed by, or reviewed by JPMorgan Chase & Co. or Anthropic PBC prior to publication.

SOURCES

References

All 47 sources below are public disclosures — earnings calls, executive interviews, Anthropic product pages, and financial/trade press. Numbered citation markers in the body ([1]–[47]) link to the matching entry here.

1. **Forbes** — How JPMorgan Chase Is Building The AI-Powered Bank Of The Future, July 1, 2026. <https://www.forbes.com/sites/bernardmarr/2026/07/01/how-jpmorgan-chase-is-building-the-ai-powered-bank-of-the-future/>
2. **Tearsheet** — JPMorgan Chase's Gen AI implementation: 450 use cases and lessons learned, Sept 17, 2025. <https://tearsheet.co/artificial-intelligence/jpmorgan-chases-gen-ai-implementation-450-use-cases-and-lessons-learned/>
3. **Klover.ai** — JPMorgan Uses AI Agents: 10 Ways to Use AI, Aug 7, 2025. <https://www.klover.ai/jpmorgan-uses-ai-agents-10-ways-to-use-ai-in-depth-analysis-2025/>
4. **Emerj** — Artificial Intelligence at JPMorgan Chase. <https://emerj.com/artificial-intelligence-at-jpmorgan-chase/>
5. **AI Magazine** — JPMorgan, OpenAI & Anthropic: Evolving Banking Operations, Oct 3, 2025. <https://aimagazine.com/news/jpmorgan-openai-anthropic-evolving-banking-operations>
6. **J.P. Morgan Payments Developer Portal** — How J.P. Morgan developers leverage AI. <https://developer.payments.jpmorgan.com/blog/guides/ai-software-development>
7. **The Digital Banker** — JPMorgan Chase's LLM Suite drives AI transformation across the enterprise, March 6, 2026. <https://thedigitalbanker.com/jpmorgan-chases-llm-suite-drives-ai-transformation-across-the-enterprise/>
8. **European Financial Review** — JPMorgan Expands AI Push with LLM Suite Integration, Oct 1, 2025. <https://www.europeanfinancialreview.com/jpmorgan-expands-ai-push-with-llm-suite-integration/>
9. **CeFPro** — JPMorgan deploys AI to help employees write performance reviews, Nov 4, 2025. <https://connect.cefpro.com/article/view/jpmorgan-deploys-ai-to-help-employees-write-performance-reviews>
10. **Fortune** — How JPMorgan's CIO is reshaping work at the bank with a \$19.8 billion annual tech and AI budget, April 29, 2026. <https://fortune.com/2026/04/29/capcom-virgin-voyages-bet-on-ai-to-reshape-gaming-and-cruise-travel/>
11. **CFO.com** — Inside Claude's rapid expansion across corporate finance, May 21, 2026. <https://www.cfo.com/news/inside-anthropic-claude-rapid-expansion-across-corporate-finance-cfo-/820806/>
12. **Fortune** — Anthropic deepens push into Wall Street with new AI agents, full Microsoft 365 integration, Moody's data partnership, May 5, 2026. <https://fortune.com/2026/05/05/anthropic-wall-street-financial-services-agents-jamie-dimon/>
13. **MindStudio** — How to Build AI Agents for Finance: Claude, OpenAI, and Anthropic's Enterprise Push, May 18, 2026. <https://www.mindstudio.ai/blog/ai-agents-for-finance-anthropic-openai-enterprise>
14. **Anthropic** — Claude for Financial Services. <https://www.anthropic.com/news/claude-for-financial-services>
15. **Pasquale Pillitteri** — Claude for Finance: Anthropic's 10 AI Agents Hit Wall Street, May 9, 2026. <https://pasqualepillitteri.it/en/news/2173/claude-finance-anthropic-10-ai-agents-2026>
16. **Constellation Research** — JPMorgan Chase, Goldman Sachs on Anthropic's Mythos, AI cyber risks, April 15, 2026. <https://www.constellationnr.com/insights/news/jpmorgan-chase-goldman-sachs-anthropics-mythos-ai-cyber-risks>
17. **Private Banker International** — JPMorgan blocks Anthropic AI models for staff in Hong Kong, June 19, 2026. <https://www.privatebankerinternational.com/news/jpmorgan-blocks-anthropic-hong-kong/>
18. **Reuters via Yahoo Finance** — JPMorgan blocks Anthropic AI access for Hong Kong staff, FT reports, June 18, 2026. <https://finance.yahoo.com/technology/ai/articles/jpmorgan-chase-cuts-off-anthropic-041524918.html>

19. **GovInfoSecurity** – JPMorgan Pulls Anthropic Claude Access in Hong Kong. <https://www.govinfosecurity.com/jpmorgan-pulls-anthropic-claude-access-in-hong-kong-a-32013>
20. **Anthropic** – Introducing the Model Context Protocol, Nov 2024. <https://www.anthropic.com/news/model-context-protocol>
21. **Model Context Protocol** – What is the Model Context Protocol (MCP)? <https://modelcontextprotocol.io/docs/getting-started/intro>
22. **Wikipedia** – Model Context Protocol. https://en.wikipedia.org/wiki/Model_Context_Protocol
23. **DEV Community** – Complete Guide to MCP (Model Context Protocol) in 2026, April 5, 2026. <https://dev.to/x4nent/complete-guide-to-mcp-model-context-protocol-in-2026-architecture-implementation-and-4a11>
24. **Kiteworks** – What Is the Model Context Protocol (MCP) - and Why It Matters for Enterprise Data Security, March 16, 2026. <https://www.kiteworks.com/cybersecurity-risk-management/model-context-protocol-enterprise-security/>
25. **Truto** – What is MCP (Model Context Protocol)? The 2026 Guide for SaaS PMs, April 3, 2026. <https://truto.one/blog/what-is-mcp-model-context-protocol-the-2026-guide-for-saas-pms/>
26. **Strategy** – Model context protocol (MCP) for enterprise AI integration. <https://www.strategy.com/software/blog/model-context-protocol-mcp-for-enterprise-ai-integration>
27. **IBM** – What is Model Context Protocol (MCP)? <https://www.ibm.com/think/topics/model-context-protocol>
28. **GuptaDeepak** – The Complete Guide to Model Context Protocol (MCP): Enterprise Adoption, Market Trends, and Implementation Strategies, Dec 2025. <https://guptadeepak.com/the-complete-guide-to-model-context-protocol-mcp-enterprise-adoption-market-trends-and-implementation-strategies/>
29. **Claude Help Center** – Set up single sign-on (SSO). <https://support.claude.com/en/articles/13132885-set-up-single-sign-on-sso>
30. **Claude Help Center** – Enable and use the Microsoft 365 connector. <https://support.claude.com/en/articles/12542951-enable-and-use-the-microsoft-365-connector>
31. **usecarly.com** – Claude + Microsoft Teams: What the Integration Can (and Can't) Do in 2026. <https://www.usecarly.com/blog/claude-microsoft-teams-integration/>
32. **UC Today** – Anthropic Expands Claude Microsoft 365 Integration for All User Plans, April 7, 2026. <https://www.uctoday.com/productivity-automation/anthropic-expands-claude-microsoft-365-integration-for-all-user-plans/>
33. **Windows Forum** – Claude Tag and the Slack Launch Threat to Microsoft Teams Copilot. <https://windowsforum.com/threads/claude-tag-and-the-slack-launch-threat-to-microsoft-teams-copilot.432750/>
34. **Anthropic** – Anthropic and Salesforce expand partnership to bring Claude to regulated industries. <https://www.anthropic.com/news/salesforce-anthropic-expanded-partnership>
35. **MindStudio** – Atlassian's MCP Server Is Now GA: How Claude Reads and Writes Jira, Confluence, and Compass via OAuth, May 4, 2026. <https://www.mindstudio.ai/blog/atlassian-mcp-server-ga-claude-reads-writes-jira-confluence-compass-oauth>
36. **pluginsforcowork.com** – Claude Cowork Connectors & Integrations (2026) - Complete List. <https://pluginsforcowork.com/guides/cowork-connectors/>
37. **howtoharder.com** – Anthropic Claude Hardening Guide. <https://howtoharder.com/guides/anthropic-claude/>
38. **platformsecurity.com** – How to Secure Your Claude Enterprise Tenant: A Settings-by-Settings Configuration Guide, April 9, 2026. <https://platformsecurity.com/blog/how-to-secure-your-claude-enterprise-tenant>
39. **Anthropic** – Agents for financial services, May 5, 2026. <https://www.anthropic.com/news/finance-agents>
40. **GitHub** – anthropics/financial-services (public plugin repository: agents, vertical skill bundles, MCP connector manifests). <https://github.com/anthropics/financial-services>
41. **J.P. Morgan Payments** – 4 Takeaways on AI in Payments from NY Tech Week 2026. <https://www.jpmorgan.com/payments/newsroom/agent-ai-payments-takeaways-2026>
42. **J.P. Morgan Asset Management** – Artificial Intelligence, market-themes research. <https://am.jpmorgan.com/us/en/asset-management/institutional/insights/market-themes/artificial-intelligence/>

- 43. **JPMorganChase** — Securing the next generation of AI agents, official technology blog, March 23, 2026. <https://www.jpmorganchase.com/about/technology/blog/securing-agentic-ai>
- 44. **Klover.ai** — JPMorgan Chase AI: Banking on AI Dominance. <https://www.klover.ai/jpmorgan-chase-ai-banking-on-ai-dominance/>
- 45. **FinAi News** — J.P. Morgan Payments not worried about escalating token cost, exec says. <https://finainews.com/payments/j-p-morgan-payments-not-worried-about-escalating-token-cost-exec-says/>
- 46. **MachineLearningMastery.com** — 5 Production Scaling Challenges for Agentic AI in 2026, March 19, 2026. <https://machinelearningmastery.com/5-production-scaling-challenges-for-agentic-ai-in-2026/>
- 47. **J.P. Morgan Payments** — Payments Outlook 2026 Trends Report Released, May 6, 2026. <https://www.jpmorgan.com/payments/newsroom/payments-outlook-2026-trends-report-released>

APPENDIX A

The Ten Claude for Financial Services Agent Templates

Referenced in Section 5.4. Each agent is a self-contained Cowork/Claude Code plugin – one system prompt plus a bundled, synced copy of the Skills it needs from its parent vertical plugin – and ships identically as a headless Claude Managed Agent template deployable via `/v1/agents` ^{[39][40]}.

METHODOLOGY NOTE

Anthropic's repository publishes the full 10-agent list with one-line functions, and publishes full skill/command tables for 5 of its 7 verticals (financial-analysis core, investment-banking, equity-research, private-equity, wealth-management) ^[40]. It does not publish an explicit per-agent “uses exactly skills [A,B,C]” table. Below, ✓ VERIFIED marks facts taken directly from Anthropic's own tables; ◆ INFERRED marks this author's mapping of an agent to its most functionally-aligned vertical skill set – a reasonable inference, not Anthropic's stated breakdown.

COVERAGE & ADVISORY

AGENT	WHAT IT DOES	SKILLS / COMMANDS	KEY MCP CONNECTORS
Pitch Agent	Comps, precedents, LBO → branded pitch deck, end to end	✓ /comps /lbo /ppt-template (core) · ◆ pitch-deck, /one-pager, /buyer-list (investment-banking)	FactSet, PitchBook, S&P Global, LSEG, Box/Egnyte
Meeting Prep Agent	Briefing pack before every client meeting	✓ competitive-analysis, deck-refresh (core) · ◆ /client-review (wealth-management)	Aiera, MT Newswires, Morningstar/FactSet, Box/Egnyte

RESEARCH & MODELING

AGENT	WHAT IT DOES	SKILLS / COMMANDS	KEY MCP CONNECTORS
Market Researcher	Sector/theme → industry overview, competitive landscape, peer comps, ideas shortlist	✓ /comps, competitive-analysis (core) · ♦ /sector, /screen, /catalysts (equity-research)	FactSet, S&P Global, Morningstar, MT Newswires, LSEG, PitchBook
Earnings Reviewer	Earnings call + filings → model update → note draft	✓ /debug-model (core) · ♦ /earnings, /earnings-preview, /model-update, /thesis (equity-research)	Aiera (call intelligence), Daloopa, FactSet, MT Newswires
Model Builder	DCF, LBO, 3-statement, comps – live in Excel	✓ /dcf /lbo /3-statement-model /comps /debug-model, xlsx-author (all core)	Daloopa, FactSet, Moody's, S&P Global

FUND ADMIN & FINANCE OPS

AGENT	WHAT IT DOES	SKILLS / COMMANDS	KEY MCP CONNECTORS
Valuation Reviewer	Ingests GP packages, runs valuation template, stages LP reporting	✓ audit-xls (core) · ♦ /returns, /portfolio (private-equity)	Chronograph, PitchBook, Box/Egnyte
GL Reconciler	Finds breaks, traces root cause, routes for sign-off	✓ GL recon, break tracing (fund-admin vertical, not itemized) + audit-xls, clean-data-xls (core)	FactSet, Morningstar, Box/Egnyte
Month-End Closer	Accruals, roll-forwards, variance commentary	✓ Accruals/roll-forwards (fund-admin vertical, not itemized) + xlsx-author (core)	FactSet, Morningstar, Box/Egnyte
Statement Auditor	Audits LP statements before distribution	✓ /debug-model (core) + NAV tie-out (fund-admin vertical)	Chronograph, Box/Egnyte

OPERATIONS & ONBOARDING

AGENT	WHAT IT DOES	SKILLS / COMMANDS	KEY MCP CONNECTORS
KYC Screener	Parses onboarding docs, runs the rules engine, flags gaps	✓ KYC document parsing, rules-grid evaluation (operations vertical, not itemized)	Box/Egnyte (document intake), Moody's (entity/credit data)

SHARED PLUMBING (VERIFIED, APPLIES TO ALL TEN)

Twelve MCP connectors live once in the financial-analysis core plugin and are shared across every agent: FactSet, Moody's, Daloopa, Morningstar, S&P Global, LSEG, PitchBook, Chronograph, Aiera, MT Newswires, Box, and Egnyte. Provider access may still require the firm's own data-license subscription – Anthropic ships the connector, not the data license ^[40]. Two deployment surfaces share one source of truth: interactive (Cowork/Code, analyst-driven) or headless (Managed Agents via /v1/agents, for nightly or multi-hour autonomous runs). The repository's own disclaimer states these agents draft work product for review – they do not execute transactions, bind risk, post to a ledger, or approve onboarding autonomously ^[40].



ARIANA.DIGITAL

AEGIS Governance Advisory

Ariana Digital LLC helps regulated enterprises — Financial Services, Healthcare, Energy, and Federal/GovCon — build defensible AI governance programs using the AEGIS Framework (Agentic Enterprise Governance and Intelligence Standard): Diagnostic, Build, and Run.

ARIANA.DIGITAL · ANTHROPIC CLAUDE PARTNER — ARIANA DIGITAL LLC